Reference Guide for the Centillion 50/100 and Model 5000BH Switches

Part No. 893-01006-C
March 1999

Bay Networks

4401 Great America Parkway
Santa Clara, CA 95052

8 Federal Street
Billerica, MA 01821

© 1999 by Bay Networks, Inc. All rights reserved.

Trademarks

Bay Networks and BN are registered trademarks of Bay Networks, Inc.

Accelar, BayStack, Centillion, Centillion 50, Centillion 100, Centillion 1000, GIGArray, SpeedView, System 5000, ATMSpeed, EtherSpeed, and TokenSpeed are trademarks of Bay Networks, Inc.

Other brand and product names are registered trademarks or trademarks of their respective holders.

Statement of Conditions

In the interest of improving internal design, operational function, and/or reliability, Bay Networks, Inc. reserves the right to make changes to the products described in this document without notice.

Bay Networks, Inc. does not assume any liability that may occur due to the use or application of the product(s) or circuit layout(s) described herein.

Portions of the code in this software product are Copyright © 1988, Regents of the University of California. All rights reserved. Redistribution and use in source and binary forms of such portions are permitted, provided that the above copyright notice and this paragraph are duplicated in all such forms and that any documentation, advertising materials, and other materials related to such distribution and use acknowledge that such portions of the software were developed by the University of California, Berkeley. The name of the University may not be used to endorse or promote products derived from such portions of the software without specific prior written permission.

SUCH PORTIONS OF THE SOFTWARE ARE PROVIDED “AS IS” AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, WITHOUT LIMITATION, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE.

In addition, the program and information contained herein are licensed only pursuant to a license agreement that contains restrictions on use and disclosure (that may incorporate by reference certain limitations and notices imposed by third parties).

RESTRICTED RIGHTS LEGEND: Use, duplication, or disclosure by the U.S. Government is subject to the restrictions set forth in subparagraph (c)(1) of the Commercial Computer Software—Restricted Rights clause of FAR52.227-19 or subparagraph (c)(1)(a) of the Rights in Technical Data and Computer Software clause of DFARS 52.227-7013, and any successor rules or regulations, whichever is applicable.

Bay Networks, Inc. Software License Agreement

NOTICE: Please carefully read this license agreement before copying or using the accompanying software or installing the hardware unit with pre-enabled software (each of which is referred to as “Software” in this Agreement). BY COPYING OR USING THE SOFTWARE, YOU ACCEPT ALL OF THE TERMS AND CONDITIONS OF THIS LICENSE AGREEMENT. THE TERMS EXPRESSED IN THIS AGREEMENT ARE THE ONLY TERMS UNDER WHICH BAY NETWORKS WILL PERMIT YOU TO USE THE SOFTWARE. If you do not accept these terms and conditions, return the product, unused and in the original shipping container, within 30 days of purchase to obtain a credit for the full purchase price.

1. License Grant. Bay Networks, Inc. (“Bay Networks”) grants the end user of the Software (“Licensee”) a personal, nonexclusive, nontransferable license: a) to use the Software either on a single computer or, if applicable, on a single authorized device identified by host ID, for which it was originally acquired; b) to copy the Software solely for backup purposes in support of authorized use of the Software; and c) to use and copy the associated user manual solely in support of authorized use of the Software by Licensee. This license applies to the Software only and does not extend to Bay Networks Agent software or other Bay Networks software products. Bay Networks Agent software or other Bay Networks software products are licensed for use under the terms of the applicable Bay Networks, Inc. Software License Agreement that accompanies such software and upon payment by the end user of the applicable license fees for such software.

2. Restrictions on use; reservation of rights. The Software and user manuals are protected under copyright laws. Bay Networks and/or its licensors retain all title and ownership in both the Software and user manuals, including any revisions made by Bay Networks or its licensors. The copyright notice must be reproduced and included with any copy of any portion of the Software or user manuals. Licensee may not modify, translate, decompile, disassemble, use for any competitive analysis, reverse engineer, distribute, or create derivative works from the Software or user manuals or any copy, in whole or in part. Except as expressly provided in this Agreement, Licensee may not copy or transfer the Software or user manuals, in whole or in part. The Software and user manuals embody Bay Networks' and its licensors' confidential and proprietary intellectual property. Licensee shall not sublicense, assign, or otherwise disclose to any third party the Software, or any information about the operation, design, performance, or implementation of the Software and user manuals that is confidential to Bay Networks and its licensors; however, Licensee may grant permission to its consultants, subcontractors, and agents to use the Software at Licensee's facility, provided they have agreed to use the Software only in accordance with the terms of this license.

3. Limited warranty. Bay Networks warrants each item of Software, as delivered by Bay Networks and properly installed and operated on Bay Networks hardware or other equipment it is originally licensed for, to function substantially as described in its accompanying user manual during its warranty period, which begins on the date Software is first shipped to Licensee. If any item of Software fails to so function during its warranty period, as the sole remedy Bay Networks will at its discretion provide a suitable fix, patch, or workaround for the problem that may be included in a future Software release. Bay Networks further warrants to Licensee that the media on which the Software is provided will be free from defects in materials and workmanship under normal use for a period of 90 days from the date Software is first shipped to Licensee. Bay Networks will replace defective media at no charge if it is returned to Bay Networks during the warranty period along with proof of the date of shipment. This warranty does not apply if the media has been damaged as a result of accident, misuse, or abuse. The Licensee assumes all responsibility for selection of the Software to achieve Licensee's intended results and for the installation, use, and results obtained from the Software. Bay Networks does not warrant a) that the functions contained in the software will meet the Licensee's requirements, b) that the Software will operate in the hardware or software combinations that the Licensee may select, c) that the operation of the Software will be uninterrupted or error free, or d) that all defects in the operation of the Software will be corrected. Bay Networks is not obligated to remedy any Software defect that cannot be reproduced with the latest Software release. These warranties do not apply to the Software if it has been (i) altered, except by Bay Networks or in accordance with its instructions; (ii) used in conjunction with another vendor's product, resulting in the defect; or (iii) damaged by improper environment, abuse, misuse, accident, or negligence. THE FOREGOING WARRANTIES AND LIMITATIONS ARE EXCLUSIVE REMEDIES AND ARE IN LIEU OF ALL OTHER WARRANTIES EXPRESS OR IMPLIED, INCLUDING WITHOUT LIMITATION ANY WARRANTY OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Licensee is responsible for the security of its own data and information and for maintaining adequate procedures apart from the Software to reconstruct lost or altered files, data, or programs.

4. Limitation of liability. IN NO EVENT WILL BAY NETWORKS OR ITS LICENSORS BE LIABLE FOR ANY COST OF SUBSTITUTE PROCUREMENT; SPECIAL, INDIRECT, INCIDENTAL, OR CONSEQUENTIAL DAMAGES; OR ANY DAMAGES RESULTING FROM INACCURATE OR LOST DATA OR LOSS OF USE OR PROFITS ARISING OUT OF OR IN CONNECTION WITH THE PERFORMANCE OF THE SOFTWARE, EVEN IF BAY NETWORKS HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. IN NO EVENT SHALL THE LIABILITY OF BAY NETWORKS RELATING TO THE SOFTWARE OR THIS AGREEMENT EXCEED THE PRICE PAID TO BAY NETWORKS FOR THE SOFTWARE LICENSE.

5. Government Licensees. This provision applies to all Software and documentation acquired directly or indirectly by or on behalf of the United States Government. The Software and documentation are commercial products, licensed on the open market at market prices, and were developed entirely at private expense and without the use of any U.S. Government funds. The license to the U.S. Government is granted only with restricted rights, and use, duplication, or disclosure by the U.S. Government is subject to the restrictions set forth in subparagraph (c)(1) of the Commercial Computer Software—Restricted Rights clause of FAR 52.227-19 and the limitations set out in this license for civilian agencies, and subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause of DFARS 252.227-7013, for agencies of the Department of Defense or their successors, whichever is applicable.

6. Use of Software in the European Community. This provision applies to all Software acquired for use within the European Community. If Licensee uses the Software within a country in the European Community, the Software Directive enacted by the Council of European Communities Directive dated 14 May, 1991, will apply to the examination of the Software to facilitate interoperability. Licensee agrees to notify Bay Networks of any such intended examination of the Software and may procure support and assistance from Bay Networks.

7. Term and termination. This license is effective until terminated; however, all of the restrictions with respect to Bay Networks' copyright in the Software and user manuals will cease being effective at the date of expiration of the Bay Networks copyright; those restrictions relating to use and disclosure of Bay Networks' confidential information shall continue in effect. Licensee may terminate this license at any time. The license will automatically terminate if Licensee fails to comply with any of the terms and conditions of the license. Upon termination for any reason, Licensee will immediately destroy or return to Bay Networks the Software, user manuals, and all copies. Bay Networks is not liable to Licensee for damages in any form solely by reason of the termination of this license.

8. Export and Re-export. Licensee agrees not to export, directly or indirectly, the Software or related technical data or information without first obtaining any required export licenses or other governmental approvals. Without limiting the foregoing, Licensee, on behalf of itself and its subsidiaries and affiliates, agrees that it will not, without first obtaining all export licenses and approvals required by the U.S. Government: (i) export, re-export, transfer, or divert any such Software or technical data, or any direct product thereof, to any country to which such exports or re-exports are restricted or embargoed under United States export control laws and regulations, or to any national or resident of such restricted or embargoed countries; or (ii) provide the Software or related technical data or information to any military end user or for any military end use, including the design, development, or production of any chemical, nuclear, or biological weapons.

9. General. If any provision of this Agreement is held to be invalid or unenforceable by a court of competent jurisdiction, the remainder of the provisions of this Agreement shall remain in full force and effect. This Agreement will be governed by the laws of the state of California.

Should you have any questions concerning this Agreement, contact Bay Networks, Inc., 4401 Great America Parkway, P.O. Box 58185, Santa Clara, California 95054-8185.

LICENSEE ACKNOWLEDGES THAT LICENSEE HAS READ THIS AGREEMENT, UNDERSTANDS IT, AND AGREES TO BE BOUND BY ITS TERMS AND CONDITIONS. LICENSEE FURTHER AGREES THAT THIS AGREEMENT IS THE ENTIRE AND EXCLUSIVE AGREEMENT BETWEEN BAY NETWORKS AND LICENSEE, WHICH SUPERSEDES ALL PRIOR ORAL AND WRITTEN AGREEMENTS AND COMMUNICATIONS BETWEEN THE PARTIES PERTAINING TO THE SUBJECT MATTER OF THIS AGREEMENT. NO DIFFERENT OR ADDITIONAL TERMS WILL BE ENFORCEABLE AGAINST BAY NETWORKS UNLESS BAY NETWORKS GIVES ITS EXPRESS WRITTEN CONSENT, INCLUDING AN EXPRESS WAIVER OF THE TERMS OF THIS AGREEMENT.

Contents

Preface

Before You Begin xvii
Related Publications xviii
How to Get Help xix

Chapter 1
Terminology and Concepts

Terminology 1-1
LAN Emulation Basics 1-1
Emulated LAN (ELAN) 1-2
LAN Emulation Client (LEC) 1-2
LAN Emulation Configuration Server (LECS) 1-3
LAN Emulation Server (LES) 1-3
Broadcast and Unknown Server (BUS) 1-3
Virtual Ports 1-4
ATM Network Interfaces and Signaling Protocols 1-6
User-to-Network Interface (UNI) 1-6
Network-to-Network Interface (NNI) 1-6
Interim Inter-Switch Signaling Protocol (IISP) 1-6
Private Network to Network Interface (PNNI) 1-7
Synchronous Optical Network (SONET) 1-9
Synchronous Digital Hierarchy (SDH) 1-9
Forwarding Protocols 1-10
Multi-Protocol Over ATM (MPOA) 1-10
Targetless ARP 1-11

Virtual LANs (VLANs) 1-11
Spanning Tree Groups 1-12
Port-Based VLANs 1-12
Protocol-Based VLANs 1-12
802.1q VLANs 1-13
Physical Ports as Members of VLANs 1-15
Multiple VLANs in a Token Ring Environment 1-16
Configuring VLANs 1-17
Virtual Ring 1-18
MultiLink Trunking 1-18
Combining MultiLink Trunking and 802.1q VLANs 1-20
Filtering 1-21
Internet Group Management Protocol/IP Multicast (IGMP/IP) 1-22
Configuration Concepts 1-26
Using the Ethernet Port of the ATM MCP 1-26
Using Token Ring Connection Modes 1-31
Switch Connection Field 1-31
Management Wrap Field 1-32
Backplane Connection Field 1-33
Front Port Connect Field 1-34
Signaling over Nonzero Virtual Ports (Associated Signaling) 1-35
Creating Redundancy with the Centillion 50/100 and Model 5000BH Switches 1-36
Spanning Tree Redundancy 1-36
PNNI and SVC/IISP Redundancy and Load Balancing 1-36
LANE Services Redundancy 1-37
Circuit Redundancy and Load Balancing 1-39
Redundant MCP 1-39
Connecting LANs Using ATM 1-40
Using ATM LAN Emulation 1-41
Using CLC Mode of ATM Switching 1-44
Using CLC Circuit Saver Switching 1-44
Using Centillion GIGArray Switching 1-46

Implementing LANE Services 1-49
Using a LECS 1-49
Using LECS Redundancy 1-50
Using the LES/BUS Directly 1-50
Using LES/BUS Redundancy 1-50
Recommendations and Tips 1-51
LAN Network Manager and Ring Parameter Server 1-52
Using LNM to Display the Switch 1-55
Configuring the Ring Number 1-56
Ring Parameter Server 1-56

Chapter 2
Configuration Examples

Token Ring Solutions 2-1
Giving a Server a Dedicated Port 2-2
Segmenting a Token Ring LAN 2-2
Replacing Bridges or Routers 2-4
Concentrating LANs into a Router 2-5
Reducing the Number of Source Route Hops 2-8
Collapsing a Token Ring Backbone 2-9
Using Distributed Backbone Switching Without LANE 2-10
Using Distributed Backbone Switching with LANE 2-12
Ethernet Solutions 2-13
Giving a Server a Dedicated Port 2-13
Segmenting an Ethernet LAN 2-13
Replacing Bridges or Routers 2-15
Concentrating LANs into a Router 2-16
Replacing a LAN Backbone 2-18

Chapter 3
MPOA in an ATM Network

MPOA Basics 3-2
MPOA Components 3-3
MPOA Clients 3-3
Ingress MPC Functionality 3-4
Egress MPC Functionality 3-5
MPOA Servers 3-5
Ingress MPS Functionality 3-6
Egress MPS Functionality 3-6
Discovery 3-6
Establishing a Shortcut 3-7
Virtual Circuit Multiplexing 3-10

Chapter 4
PNNI in an ATM Network

What Is PNNI 4-1
Using PNNI in the ATM Network 4-2
Using PNNI with the Centillion 50/100 or Model 5000BH Switches 4-3
Using PNNI with IISP 4-9
Configuring PNNI 4-14
Addressing Schemes 4-14
Configuring PNNI with SpeedView 4-17
PNNI Features 4-18
Logical Links 4-19
Multilevel Designated Transit Lists 4-21
Border Nodes 4-22
Address Summarization 4-22
Call Balancing in PNNI Networks 4-25

Chapter 5
IISP in an ATM Network

Where IISP Fits in an ATM Network 5-2
IISP Link Configuration 5-3
IISP Route Table 5-3
Route Entry Table 5-5
Link Groups and Loop Prevention 5-6
How a Link Group Provides Loop Prevention 5-7
Example of IISP and Redundant Links in a Network 5-11
Example of Link Groups in a Network 5-13

Chapter 6
Packet Filter Applications

Traffic Management in Switched Internetworks 6-2
Introduction to Packet Filtering 6-3
Packet Filter Examples 6-6
Enhanced All Routes Explorer Reduction in Meshed Networks 6-7
General Explorer Filters for Token Ring Networks 6-9
Dropping All Routes and Spanning Tree Explorers 6-9
Dropping Spanning Tree Explorers 6-9
NetBIOS Filters 6-10
Dropping NetBIOS Frames (802.3 Format) from Ethernet Ports 6-10
Dropping NetBIOS Frames from Token Ring Ports 6-11
Dropping Single Route and All Routes NetBIOS Explorers from Token Ring Ports 6-11
SNA Filters 6-12
Forwarding SNA Frames and Dropping All Other Frames from Token Ring Ports 6-12
Forwarding SNA Frames (802.3) and Dropping All Other Frames from Ethernet Ports 6-12
Forwarding SNA Frames (Ethernet V2.0) and Dropping All Other Frames from Ethernet Ports 6-13
IPX Filters 6-13
Dropping Token Ring IPX Traffic and Forwarding All Others 6-13
Dropping IPX Frames and Forwarding All Others on Ethernet Ports 6-14

IP Filters 6-15
Dropping IP Frames (Not Including ARP and RARP) from Token Ring Ports 6-15
Dropping Version 2.0 IP Frames (Not Including ARP and RARP) from Ethernet Ports 6-15
Dropping ARP and RARP Frames from Token Ring Ports 6-16
Dropping ARP and RARP Frames from Ethernet Ports 6-16
Other Protocol Filters 6-17
Dropping Banyan Frames from Token Ring Ports 6-17
Dropping DEC LAT Frames from Ethernet Ports 6-17
Monitor and Additional Destination Filters 6-18
Monitoring All Frames on Token Ring Ports 6-18
Monitoring Frames on Ethernet Ports 6-19
Forwarding Frames to Specific Token Ring Ports 6-19
Advanced Filters for Token Ring 6-20
Restricting IPX Get Nearest Server Broadcasts to the Local Ring 6-22
Restricting IPX Get Nearest Server Broadcasts to Workgroup Rings 6-23
Restricting RIP and SAP Broadcasts to the Local Ring 6-24
Restricting RIP and SAP Broadcasts to within the Rings for a Workgroup 6-24
Reducing HP Laser Jet Printer Get General File Server Broadcasts 6-25
Restricting HP Laser Jet Printer Get General File Server Broadcasts to a Workgroup 6-25
Outbound Filters 6-26
Outbound Filter Examples 6-26
Common SAPs 6-27
Ethernet Protocol Types 6-28

Chapter 7
NetBIOS Protocol

Name Filters and Examples 7-3

Appendix A
Switch Module and Software Matrix

Switch Modules and Minimum Software Requirements A-1
Switch Image Compatibility A-5

Index

Figures

Figure 1-1. Soft PVCs in a PVC/SVC Network 1-9
Figure 1-2. Multiple VLANs in a Single Spanning Tree Group 1-12
Figure 1-3. Spanning Tree Groups with Multiple VLANs 1-13
Figure 1-4. Example of 802.1q Tagged Ethernet Frame 1-14
Figure 1-5. 802.1q Tagged VLANs 1-15
Figure 1-6. MCP with Multiple VLANs 1-17
Figure 1-7. Example of MultiLink Trunking 1-19
Figure 1-8. MultiLink Trunking and 802.1q in a Network 1-20
Figure 1-9. Application Behavior and Characteristics 1-24
Figure 1-10. Secondary-Path Management Information Request: Out of Band 1-27
Figure 1-11. Primary-Path Management Information Request: In Band 1-28
Figure 1-12. Dual-Path Management Information Request: Not Recommended 1-28
Figure 1-13. Dual-Path Management Information Request: Recommended Configuration 1-30
Figure 1-14. Packets Switched with the Switch Connect Option 1-31
Figure 1-15. Packets Wrapped Using the Wrapped Option 1-32
Figure 1-16. Model 5000BH Switch with a Split Backplane 1-33
Figure 1-17. Front Port Field Options 1-34
Figure 1-18. Signaling over Nonzero Virtual Ports 1-35
Figure 1-19. Redundant and Cooperating LES/BUS Pairs 1-38
Figure 1-20. The Switch in a Centralized ATM Switch Configuration 1-42
Figure 1-21. The Switch in a Distributed ATM Switch Configuration 1-43
Figure 1-22. Single Circuit Saver LAN (Token Ring Example) 1-45
Figure 1-23. Multiple Circuit Saver LANs (Ethernet Example) 1-45
Figure 1-24. GIGArrays Connected Using Circuit Saver Bridging 1-48
Figure 1-25. Internal and External Rings Connected by Bridge Ports on a Switch 1-53
Figure 1-26. All External Rings Connected to One Internal Ring 1-54
Figure 1-27. Each Bridge Group Displayed as a Separate Internal Ring 1-54

Figure 2-1. Segmenting a Token Ring Using a Concentrator Port 2-3
Figure 2-2. Segmenting a Token Ring Using the RI or RO Port 2-3
Figure 2-3. Replacing Source Route Bridges or Routers 2-4
Figure 2-4. Token Ring Network with No Growth Capacity 2-5
Figure 2-5. Switch Concentrating Multiple Rings as a Single Bridged Subnet 2-6
Figure 2-6. Switch Concentrating Multiple Rings as Multiple Routed Subnets 2-7
Figure 2-7. Reducing the Number of Source Route Hops 2-8
Figure 2-8. Collapsing All Ring Segments to a High-Density Switch 2-9
Figure 2-9. Distributed Backbone Switching Without LANE 2-10
Figure 2-10. Distributed Backbone Switching with LANE 2-12
Figure 2-11. Segmenting an Ethernet LAN 2-14
Figure 2-12. Replacing an Ethernet Bridge 2-15
Figure 2-13. Ethernet Network with No Growth Capacity 2-16
Figure 2-14. Switch Concentrating Multiple Segments as a Single Bridged Subnet .. 2-17
Figure 2-15. Ethernet Backbone Using Routers and Bridges 2-18
Figure 2-16. Increasing Ethernet Backbone Throughput with a Centillion 100 Switch 2-19

Figure 3-1. MPOA Server and Components 3-5
Figure 3-2. Example of Data Pack Flow with and Without MPOA 3-8
Figure 3-3. Establishing a Shortcut 3-8
Figure 3-4. Internal MPOA Shortcut 3-10
Figure 3-5. Multiplexing Virtual Circuits 3-10

Figure 4-1. PNNI Network Using Centillion 50/100 and Model 5000BH Switches 4-4
Figure 4-2. Good and Inadequate PNNI Network Designs 4-6
Figure 4-3. Sample Core Designs 4-7
Figure 4-4. ATM Network with PNNI Core and IISP Edge 4-11
Figure 4-5. Example of Inadequate IISP/PNNI Network Designs 4-13
Figure 4-6. PNNI Dialog Box 4-17
Figure 4-7. Logical vs. Physical Links 4-20
Figure 4-8. Multilevel Hierarchy of Centillion 50/100 Switches 4-21
Figure 4-9. Peer Group with Five Border Nodes 4-22
Figure 4-10. Address Summarization 4-24
Figure 4-11. PNNI Call Balancing 4-25

Figure 5-1. IISP Between Private ATM Network Switches 5-2
Figure 5-2. Configure Call Route Dialog Box with Route Entry Table 5-4
Figure 5-3. Example of Link Groups and Loop Prevention 5-7
Figure 5-4. Example of IISP Connectivity and Redundant Links 5-11
Figure 5-5. Link Groups in an IISP Network 5-13
Figure 5-6. Break in the IISP Network 5-16

Figure 6-1. Token Ring Packet Filter Offsets 6-3
Figure 6-2. Information Field for an LLC Frame 6-4
Figure 6-3. Ethernet Frame Format Showing Start of Filter Offsets 6-4
Figure 6-4. Filter That Drops All Fs Broadcast with DSAP and SSAP of E0 6-5
Figure 6-5. Using the CN_BN Filters with Redundant Token Ring Switches 6-8
Figure 6-6. Filter to Drop All Routes and Spanning Tree Explorers 6-9
Figure 6-7. Filter to Drop Spanning Tree Explorers 6-9
Figure 6-8. Filter to Drop NetBIOS Frames (802.3) from an Ethernet Port 6-10
Figure 6-9. Filter to Drop NetBIOS Frames from a Token Ring Port 6-11
Figure 6-10. Filter to Drop Single and All Routes NetBIOS Explorers 6-11
Figure 6-11. Filter to Forward Only SNA Frames on Token Ring Ports 6-12
Figure 6-12. Filter to Forward Only SNA Frames (802.3) on Ethernet Ports 6-12
Figure 6-13. Filter to Forward Only SNA Frames (2.0) on Ethernet Ports 6-13
Figure 6-14. Filter to Drop IPX Traffic on Token Ring Ports 6-13
Figure 6-15. Filter to Drop IPX Frames (V2.0) on Ethernet Ports 6-14
Figure 6-16. Filter to Drop All IP Frames Except ARP and RARP from Token Ring Ports 6-15
Figure 6-17. Filter to Drop All IP Frames (V2.0) Except ARP and RARP from Ethernet Ports 6-15
Figure 6-18. Filter to Drop ARP and RARP Frames from Token Ring Ports 6-16
Figure 6-19. Filter to Drop ARP and RARP Frames from Ethernet Ports 6-16
Figure 6-20. Filter to Drop Banyan Frames from Token Ring Ports 6-17
Figure 6-21. Filter to Drop DEC LAT Frames from Ethernet Ports 6-17
Figure 6-22. Filter to Send Frames on Token Ring Ports to a Monitor Destination 6-18
Figure 6-23. Filter to Send Frames from Ethernet Ports to a Monitor Destination 6-19
Figure 6-24. Filter to Forward Frames to Specific Token Ring Ports 6-19
Figure 6-25. Directing Routed Protocols to Router Ports and Switching Layer 2 Locally 6-20
Figure 6-26. Filters to Send Frames on Ethernet Ports to a Monitor Destination 6-21

Figure 6-27. Filters to Restrict IPX GNS Broadcasts to the Local Ring 6-22
Figure 6-28. Filter to Restrict IPX GNS Broadcasts to Workgroup Rings 6-23
Figure 6-29. Filter to Restrict RIP and SAP Broadcasts to the Local Ring 6-24
Figure 6-30. Filter to Restrict RIP and SAP Broadcasts to Within the Rings for a Workgroup 6-24
Figure 6-31. Filter to Reduce HP Laser Jet Printer Get General File Server Broadcasts 6-25
Figure 6-32. Filter to Restrict HP Laser Jet Get General File Server Broadcasts to a Workgroup 6-25
Figure 6-33. Filter to Copy Outgoing IP Packets to a Monitor Destination 6-26
Figure 6-34. Filter to Drop All Outgoing UDP Packets from Ethernet Ports 6-26
Figure 6-35. Filter to Drop All Outgoing PING (ICMP) Packets from Ethernet Ports .. 6-27

Figure 7-1. NetBIOS Filter Assignments 7-4

Tables

Table 1-1. Port and Ring Assignments on a Model 5000BH Switch 1-34
Table 1-2. CLC Circuit Saver versus Turbo 1-44
Table 1-3. ATM MCP Memory Requirements 1-52

Table 3-1. Possible SVC flows 3-11

Table 4-1. Addressing Scheme for a PNNI Network 4-4
Table 4-2. IISP Route Table in the Data Center Edge Switches (Switch A) 4-12
Table 4-3. IISP Route Table in the Wiring Closet (Switch 1) 4-12
Table 4-4. Link Group Table 4-12
Table 4-5. Registered NSAP ATM Address Fields 4-14
Table 4-6. Sample User-Defined NSAP Fields 4-15

Table 5-1. Centillion 100 Switch 2 Routing Table 5-8
Table 5-2. Switch 2 Link Groups 5-8
Table 5-3. Model 5000BH Switch Routing Table 5-9
Table 5-4. IISP Link Groups 5-10
Table 5-5. Centillion 100 Switch 1 Routing Table 5-12
Table 5-6. Model 5000BH Switch Routing Table 5-12
Table 5-7. Centillion 100 Switch 2 Routing Table 5-12
Table 5-8. Switch 1 (39.01) IISP Routing Table 5-13
Table 5-9. Switch 1 (39.01) Link Groups 5-14
Table 5-10. Switch 2 (39.02) IISP Routing Table 5-14
Table 5-11. Switch 2 (39.02) Link Groups 5-14
Table 5-12. Switch 3 (39.03) Routing Table 5-14
Table 5-13. Switch 3 (39.03) Link Groups 5-14
Table 5-14. Switch 4 (39.04) IISP Route Table 5-15
Table 5-15. Switch 4 (39.04) Link Groups 5-15

Table 6-1. Common SAPs 6-27
Table 6-2. Ethernet Protocol Types 6-28

Table 7-1. NetBIOS Features and Frame Types 7-2
Table 7-2. NetBIOS Sample Data 7-5

Table A-1. Centillion 50/100 Switch Module and Minimum Required Software Matrix A-1
Table A-2. Model 5000BH Switch Modules A-3
Table A-3. Switch Image Compatibility Matrix A-5

Preface

Congratulations on your purchase of SpeedView™ 4.0 for Windows and the Centillion 50™ or Centillion 100™ or Model 5000BH switch. SpeedView is the graphical user interface (GUI) application that allows you to configure and monitor the Centillion 50/100 and Model 5000BH switches and the network traffic.

Use the concepts and reference materials in this guide to help configure your network for optimum performance.

Note: In this guide, the Centillion 50/100 and Model 5000BH switches are referred to as the switch.

Before You Begin

This guide provides configuration examples and reference information for the following subjects:

- Terminology and concepts
- Configuration examples
- MPOA in the ATM network
- PNNI in the ATM network
- IISP in the ATM network
- Packet filter applications
- NetBIOS Protocol
- Switch modules and software compatibility

Separate guides contain specific information about installing and using the Centillion 50/100 and Model 5000BH switches. See "Related Publications" for a list of documents that supplement this guide.

Related Publications

For more information about SpeedView and the Centillion 50/100 and Model 5000BH switches, refer to the following Bay Networks® publications:

- *Using SpeedView 4.0 for Windows* (Bay Networks part number 893-891-D)
- *Reference Guide for the Centillion Command Line Interface* (Bay Networks part number 893-00985-C)

The following publications contain additional information about the concepts described in this book:

- *ATM User-Network Interface (UNI) Specification, Version 3.1*
- *ATM User-Network Interface (UNI) Specification, Version 3.0*
- *LAN Emulation Over ATM, Version 1.0*
- *LAN Emulation User To Network Interface (LUNI), Version 2.0*
- *Interim Inter-switch Signaling Protocol (IISP) Specification, Version 1.0*
- *Private Network-to-Network Interface Specification, Version 1.0 (PNNI 1.0)*
- *Multi-Protocol Over ATM, Version 1.0 (AF-MPOA-0087.000)*

You can now print Bay Networks technical manuals and release notes free, directly from the Internet. Go to *support.baynetworks.com/library/tpubs/*. Find the Bay Networks product for which you need documentation. Then locate the specific category and model or version for your hardware or software product. Using Adobe Acrobat Reader, you can open the manuals and release notes, search for the sections you need, and print them on most standard printers. You can download Acrobat Reader free from the Adobe Systems Web site, *www.adobe.com*.

You can purchase Bay Networks documentation sets, CDs, and selected technical publications through the Bay Networks Collateral Catalog. The catalog is located on the World Wide Web at *support.baynetworks.com/catalog.html* and is divided into sections arranged alphabetically:

- The “CD ROMs” section lists available CDs.
- The “Guides/Books” section lists books on technical topics.
- The “Technical Manuals” section lists available printed documentation sets.

Make a note of the part numbers and prices of the items that you want to order. Use the “Marketing Collateral Catalog description” link to place an order and to print the order form.

How to Get Help

For product assistance, support contracts, or information about educational services, go to the following URL:

http://www.baynetworks.com/corporate/contacts/

Or telephone the Bay Networks Technical Solutions Center at:

800-2LANWAN

Chapter 1
Terminology and Concepts

This chapter presents the terminology related to the features that are supported by SpeedView, the Centillion 50/100 and Model 5000BH switches, and LAN emulation, as well as concepts and examples of how the Centillion 50/100 and Model 5000BH switches work within a network. The following topics are discussed in this chapter:

- Terminology (this page)
- Configuration concepts (page 1-26)
- Implementing LANE services (page 1-49)
- LAN network manager and ring parameter server (page 1-52)

Terminology

The following section introduces the concepts and terminology used in this guide including LAN emulation terms, signaling protocols, and port types.

LAN Emulation Basics

This section defines ATM LAN emulation terminology as it applies to the Centillion and Model 5000BH switch platforms.

LAN emulation (LANE) refers to the services and protocols defined in the ATM Forum LAN Emulation Version 1.0 and 2.0 specifications. LANE Version 2.0 is divided into two specifications: LANE User-to-Network Interface (LUNI 2.0) and LANE Network-to-Network Interface (LNNI 2.0).

LANE allows token ring and Ethernet clients on LAN-to-ATM bridge devices to communicate with direct-attached ATM servers and with other LAN clients on other LAN-to-ATM bridges transparently across ATM. Communication is accomplished by creating token ring or Ethernet emulated LANs (ELANs) in the ATM network and then configuring ATM servers and LAN-to-ATM switches to connect to them. In the case of the Centillion 50/100 and Model 5000BH switches, the user can configure LAN emulation client (LEC) virtual ports to participate in ELANs.

LUNI 2.0 is compatible with LANE 1.0 and provides several new features including VCC multiplexing, MPOA, and other enhancements to the communication between LANE clients and LANE services (LES/BUS and LECS). The Centillion 50/100 and Model 5000BH switches support the LUNI 2.0 client.

For more information about LANE, refer to the ATM Forum specifications *LAN Emulation Over ATM, Version 1.0.* and *LAN Emulation Over ATM, Version 2.0* (LUNI and LNNI).

Emulated LAN (ELAN)

Emulated LAN refers to an implementation of a virtual LAN using the ATM Forum's LAN Emulation Specification. A Centillion 50/100 or Model 50000BH ATM switch network using LANE-based virtual ports (LECs), third party ATM host adapter cards, and ATM routers that support LANE clients may participate in an emulated LAN.

LAN Emulation Client (LEC)

A LAN emulation client performs data forwarding over ATM virtual clients, address resolution, and other control functions. The LEC provides a MAC-level emulated Ethernet IEEE 802.3 or token ring IEEE 802.5 service interface to higher-level software. The LEC implements an ATM Forum LANE standards-compliant LUNI when communicating with other LECs in an ELAN.

For more information about LEC, refer to the ATM Forum specifications *LAN Emulation Over ATM, Version 1.0* and *LAN Emulation Over ATM, Version 2.0* (LUNI and LNNI).

LAN Emulation Configuration Server (LECS)

The LAN emulation configuration server assigns individual LECs to different ELANs based upon its policies, configuration database, and information provided by LECs. The LECS gives the client the ATM address of the LAN emulation server (LES) providing services for that specific ELAN.

For more information about LECS, refer to the ATM Forum specifications *LAN Emulation Over ATM, Version 1.0* and *LAN Emulation Over ATM, Version 2.0* (LUNI and LNNI).

LAN Emulation Server (LES)

The LAN emulation server provides the control coordination function for the ELAN. The LES registers and resolves MAC addresses and route descriptors to ATM addresses. The LES may register its LAN destinations with the LECS. A LAN emulation client queries the LES when the client wishes to resolve a MAC address or route descriptor to an ATM address. The LES either responds directly to the LEC using the Smart LES option or forwards the query to other LECs so that they may respond.

For more information about LES, refer to the ATM Forum specifications *LAN Emulation Over ATM, Version 1.0* and *LAN Emulation Over ATM, Version 2.0* (LUNI and LNNI).

Broadcast and Unknown Server (BUS)

The broadcast and unknown server handles data sent by a LEC to the broadcast MAC address, all multicast traffic, and initial unicast frames that are sent by a LEC before the data direct target ATM address has been resolved.

For more information about BUS, refer to the ATM Forum specifications *LAN Emulation Over ATM, Version 1.0* and *LAN Emulation Over ATM, Version 2.0* (LUNI and LNNI).

Virtual Ports

A virtual port (Vport) represents a collection of ATM VCs that connect a physical port to ports on remote switches. There are two types of virtual ports: Centillion LAN client (CLC) and LAN emulation client (LEC). CLC Vports rely on user-defined permanent virtual circuits (PVCs) and permanent virtual paths (PVPs) to connect to other CLC Vports in remote switches across an ATM infrastructure. LEC virtual ports rely on switched virtual circuits (SVCs) and the IISP or PNNI standards for interswitch communications. The operation of these virtual ports is compliant with the ATM Forum LANE specification versions 1.0 and 2.0 for LAN emulation clients.

Selecting either Vport type—CLC or LEC—depends on your network and your current needs. Because a CLC uses PVCs, clients use static paths to communicate. For small networks, this simplicity optimizes speed, robustness, and efficiency of the network. The network is highly fault tolerant because it is fully distributed with no single point of failure.

For medium to large networks, a LEC is simpler to configure because the necessary circuits are automatically established. A LEC also allows standards-compliant interoperability with ATM routers and ATM servers.

You can configure CLC or LEC Vports to forward packets based on one of two forwarding methods: Circuit Saver mode or Turbo mode. For more information about forwarding methods, see "Turbo Forwarding" and "Circuit Saver Forwarding" on page 1-4.

Turbo Forwarding

Turbo mode maximizes throughput and is ideal for network environments requiring optimal performance. However, Turbo mode also uses a large number of VCs in your network. You must consider VC usage when deciding how to scale the number of physical switches in a domain.

Circuit Saver Forwarding

In some environments, you may need to conserve your virtual paths and/or virtual circuits. Circuit Saver forwarding optimizes ATM circuit resources by using a minimum set of VCs to establish communication. Circuit saver uses fewer virtual circuits for the communication between local LAN ports and remote switches. When a CLC is configured for Circuit Saver, the user creates these virtual circuits

by configuring PVCs to connect to remote devices. When a LEC is configured for Circuit Saver, the LEC automatically creates a single SVC between the VLAN and each remote LEC as necessary. Because Circuit Saver uses fewer VCs, you can scale to a larger network.

When you plan your network consider that the VC count is only relevant in networks that have more than 30 switches and that Turbo mode provides higher performance than Circuit Saver mode.

Centillion LAN Client (CLC) Virtual Ports

When using a CLC Circuit Saver mode or CLC Turbo mode (GIGArray™), each individual member of the LAN group must have circuits defined with every other member. All Centillion 50/100 and Model 5000BH switch virtual ports in a CLC Circuit Saver LAN or GIGArray LAN should have a complete mesh of configured PVCs or PVPs, respectively.

Connections in a CLC Circuit Saver LAN are configured using permanent virtual circuits (PVCs). Nonzero virtual path identifiers (VPIs) are supported. The virtual channel identifier (VCI) should start at 32.

Connections in a CLC GIGArray LAN are configured using PVPs. The virtual paths must have a nonzero VPI between 1 and 255. The starting VCI value within the PVP can be configured for specialized applications.

LEC Virtual Ports

LEC Circuit Saver Vports are ideal for multi-vendor LANE interoperability where you need to optimize SVC usage.

A Turbo LEC Vport is a high-performance virtual port used for LAN emulation. The switch-to-switch throughput performance for a switch has been demonstrated in excess of 700,000 packets per second (pps). Like Turbo CLCs, using Turbo LEC increases the SVCs used by the switch.

Note: Only Turbo Ethernet and Transparent Bridging mode token ring ELANs are currently supported. Source route bridging (SRB) and source route transparent (SRT) bridging mode token ring LANE is not supported in Turbo mode.

ATM Network Interfaces and Signaling Protocols

This section introduces the types of ATM network interfaces and the signaling protocols used as ATM network interfaces in a Centillion 50/100 and Model 5000BH switch environment.

User-to-Network Interface (UNI)

The user-to-network interface is the interface between an end point of an ATM network and the switch. It represents the user connection to a Centillion 50/100 or Model 5000BH switch. The UNI can be the interface between an end point and a switch or between a switch and a router.

For more information about UNI 3.0 and UNI 3.1, refer to *ATM User-Network Interface (UNI) Specification, Version 3.0* and *ATM User-Network Interface (UNI) Specification, Version 3.1*.

Network-to-Network Interface (NNI)

The network-to-network interface is the interface between two switches. Private Network to Network Interface (PNNI) is a call routing protocol used between private network switches located within a building or campus environment.

Interim Inter-Switch Signaling Protocol (IISP)

IISP provides an interface between two ATM switches for forwarding call setup messages, allowing switched virtual circuits (SVCs) to be set up across them.

IISP, as defined by the ATM Forum, is an interim protocol for the private NNI. IISP is built upon the ATM Forum UNI 3.1 specification, with optional support for the ATM Forum UNI 3.0 protocol.

The call routing tables for each link must be manually configured to support a default path, alternate path routing for failure protection, and load balancing.

For more information about IISP, refer to *Interim Inter-switch Signaling Protocol Specification, Version 1.0*. For a discussion about using IISP in an ATM network, see Chapter 5, “IISP in an ATM Network.”

Private Network to Network Interface (PNNI)

PNNI is a routing protocol for distributing topology information between private ATM switches and clusters of private ATM switches, with a second protocol defined for signaling. PNNI is also sometimes defined as Private Network Node Interface.

The PNNI 1.0 protocol is defined by the ATM Forum to be compatible with the ATM Forum UNI 3.1 specification, supporting all UNI 3.1 and some UNI 4.0 capabilities.

The PNNI topology protocol includes a hierarchy mechanism with the ability to automatically configure itself in large ATM networks in which the address structure reflects the topology. PNNI topology and routing are based on a link-state routing protocol.

PNNI signaling—message flows used to establish point-to-point and point-to-multipoint ATM connections—is based upon the ATM Forum UNI signaling, with added mechanisms to support source routing, crankback, and alternate routing in case of connection setup failure.

For a discussion about using PNNI in an ATM network, see Chapter 4, “PNNI in an ATM Network.”

Multilevel Designated Transit List

The Centillion 50/100 switch supports multilevel designated transit lists (DTLs). The DTL is a stack of designated transit list information elements including all logical nodes (links) to a destination that are known to the originator, as well as port information about the destination. A multilevel DTL includes peer groups in different levels of the PNNI hierarchy. This means that two Centillion 50/100 and Model 5000BH switches in two separate peer groups can communicate over PNNI.

For more discussion about multilevel DTL and PNNI, see Chapter 4, “PNNI in an ATM Network.”

Note: The Model 5000BH switch can participate in a multilevel PNNI hierarchy but cannot act as the peer group leader. You must use a Centillion 1000™ switch or a third party switch to provide peer group leader functionality.

Soft Permanent Virtual Circuits

Soft PVCs (SPVCs) allow you to connect Centillion PVC end points (using the same PVC encapsulation method) across a PNNI-routed ATM network without having to configure PVCs at every intermediate switch in the path. With SPVCs, you retain the dynamic nature of PNNI, along with PNNI's path selecting and rerouting features by connecting the PVC end stations with a virtual SVC connection.

Although many large networks load balance multiple ATM links between switches, some customers still extend their networks over a single link to a wiring closet or a remote building within the campus. If LANE is not in use, the customer must build a network of PVCs to interconnect the switches, ATM servers, and routers. Soft PVCs allow the customer to create a link to the ingress ATM switch port across multiple switches to the remote egress ATM switch.

Figure 1-1 shows a Centillion 50 edge switch with a PVC to a Centillion 100 switch that is part of a PNNI network. The Centillion 50 edge switch is using the core PNNI network as an SVC to the other Centillion 50 edge switch. The SPVC provides PNNI resiliency to PVC connections between edge switches.

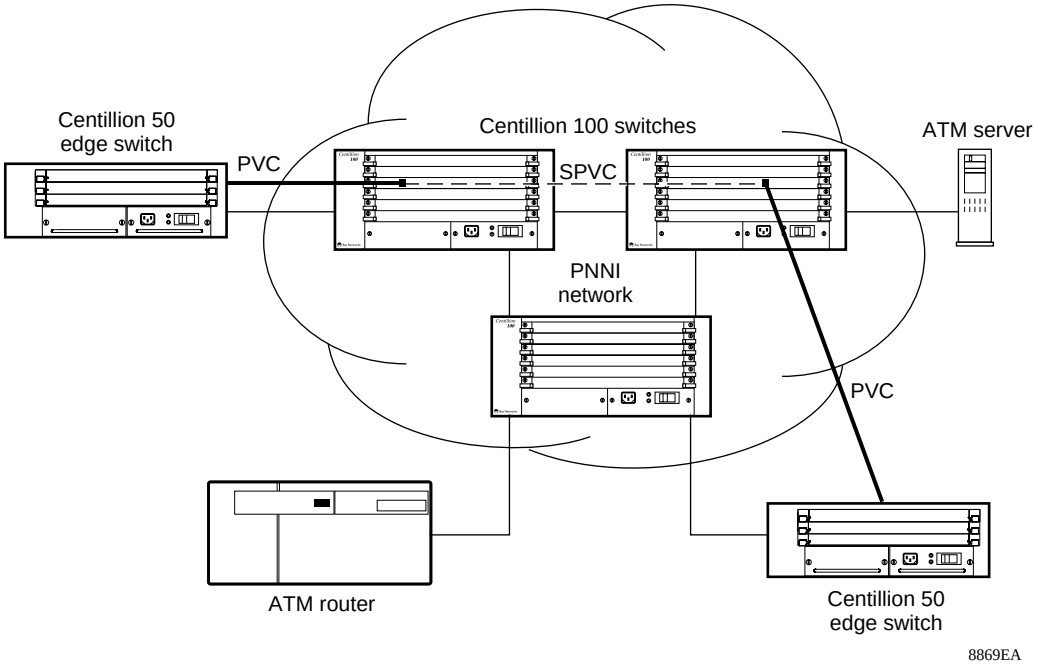

Figure 1-1. Soft PVCs in a PVC/SVC Network

Synchronous Optical Network (SONET)

SONET is a family of fiber optic transmission rates ranging from 51.84 megabits per second (Mb/s) to 13.22 gigabits per second (Gb/s). SONET offers the flexibility to transport many digital signals with different capacities and to provide a design standard for manufacturers.

Synchronous Digital Hierarchy (SDH)

SDH is the international standard for optical fiber-based synchronous data transmission. SDH is the European equivalent of SONET.

Forwarding Protocols

This section describes the forwarding protocols that the Centillion and Model 5000BH switches support.

Multi-Protocol Over ATM (MPOA)

MPOA is the ATM forum standard that provides distributed Layer 3 forwarding capabilities for LAN traffic across an ATM network without requiring the traffic to pass through a traditional multiprotocol router.

For a general discussion about MPOA, see Chapter 3, "MPOA in an ATM Network." For more information about MPOA, refer to The ATM Forum specification *Multi-Protocol Over ATM, Version 1.0.*

MPOA Client (MPC)

The MPC is the MPOA client software process that operates in an ATM edge device or LANE host that monitors Layer 3 traffic flows. The MPC performs Layer 3 forwarding, but does not run internetwork layer routing protocols.

MPOA Server (MPS)

The MPS is the MPOA server software process that communicates with a router for layer 3 information. The MPS also communicates with LANE processes for layer 2 information. Each MPS in a network communicates with other MPOA devices to resolve layer 3 addresses to ATM addresses.

Next Hop Resolution Protocol (NHRP)

NHRP is an address resolution protocol used by a next hop client (NHC) or a next hop server (NHS) to establish SVCs across subnet boundaries for MPOA.

Targetless ARP

Targetless Address Resolution Protocol (ARP) is used to expedite address learning in a dynamic environment where end stations are moved on a frequent basis and seen as a roaming MAC address.

When a station moves to a new port, the master control processor (MCP) of the receiving switch broadcasts a targetless ARP that contains the new location of the end station, or MAC address, to its emulated LAN (ELAN). The switch that owned the station previously will delete its MAC address entry and create a new entry in its forwarding database (FDB) or bridge table that indicates the new location of the MAC address. All of the switches that had the MAC address in their FDBs will update their respective FDBs with the new location of the MAC address.

Targetless ARP is limited to transparent bridging (TB). It is not supported for source route bridging (SRB) or source route transparent bridging (SRT) and addresses station movement within a single ELAN.

Virtual LANs (VLANs)

A virtual LAN is a set of components connected together in a single layer 2 switching domain (also referred to as a bridging or broadcast domain). Layer 3 routers or bridges provide connectivity between different VLANs.

Centillion 50/100 and Model 5000BH switches support three types of virtual LANs: port-based VLANs, protocol-based VLANs, and 802.1q tagged VLANs. For more information about VLAN types, see “Port-Based VLANs,” “Protocol-Based VLANs,” and “802.1q VLANs.”

A single VLAN may be thought of as a multiport bridge, some ports of which can be physically seen on the switch and some of which are virtual and defined through configuration. VLANs operate under the following rules:

- A single switch may have multiple VLANs configured.
- Each VLAN within a switch may be an independent spanning tree group, or it may share a spanning tree group with other VLANs.

VLANs can include physical and virtual ports. Physical ports are LAN ports such as Ethernet, Fast Ethernet, and token ring. The only port combination not allowed in the same VLAN is Ethernet and token ring.

Spanning Tree Groups

Centillion 50/100 and Model 5000BH switches support the Spanning Tree Protocol. In addition, a switch can support multiple spanning tree instances within the same box. The collection of physical and virtual ports considered part of one spanning tree instance is called a spanning tree group (STG). The switch supports multiple spanning tree groups, and multiple VLANs can share a spanning tree group. A VLAN can belong to one spanning tree group only.

Figure 1-2 illustrates two spanning tree groups with multiple VLANs in each.

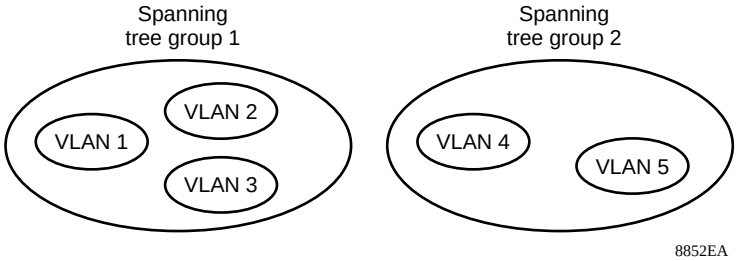

Figure 1-2. Multiple VLANs in a Single Spanning Tree Group

Each physical (token ring or Ethernet) or virtual port can belong to one and only one spanning tree group. Trunk ports can belong to more than one spanning tree group.

Port-Based VLANs

Users define port-based VLANs by grouping physical and virtual ports together. All traffic that does not match the protocol type or VLAN ID of other protocol-sensitive or 802.1q VLANs is forwarded to a port-based VLAN.

Protocol-Based VLANs

Protocol-based VLANs classify incoming packets based on the protocol type field in the packet header and the ingress port number. A port can belong to multiple protocol-based VLANs, but only if the VLANs use different protocols.

In Figure 1-3, ports 1 through 4 are in spanning tree group 1. VLANs IP1, IP2, P1, P2, and IPX split spanning tree group 1 into five separate broadcast domains. Ports 5 through 8 comprise spanning tree group 2, which is shared by one port-based VLAN (P3) and one protocol-based VLAN (IP3).

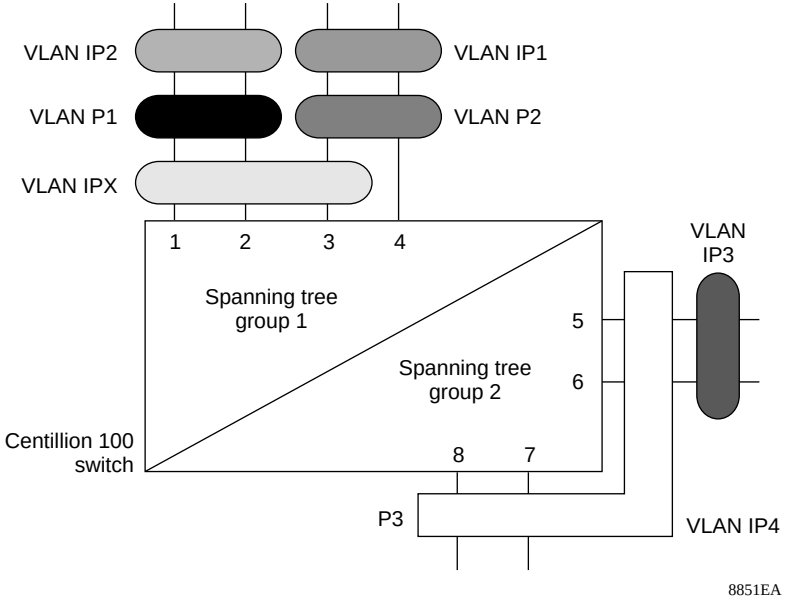

Figure 1-3. Spanning Tree Groups with Multiple VLANs

Port 1 forwards IP packets to VLAN IP2 and IPX packets to VLAN IPX. All other packets are forwarded to port-based VLAN P1. Similarly, Port 4 forwards all IP traffic to VLAN IP1 and all other traffic to port-based VLAN P2.

802.1q VLANs

The Centillion 50/100 and Model 5000BH switches support the IEEE 802.1q specification for “tagging” frames. The specification defines a method for coordinating VLANs across multiple switches. In the specification, an additional 4-octet (“tag”) header is inserted in a frame after the source MAC address and before the frame type as shown in Figure 1-4. By coordinating VLAN IDs across multiple switches, VLANs can be extended to multiple switches.

The tag has two primary components, the Tag Protocol Identifier (TPID) and the Tag Control Information (TCI). Each component is two bytes long. The TPID field identifies the packed as a tagged frame. Its value is unique with respect to Ethernet-encapsulated protocol types. The TCI field contains the following three components:

- The User Priority field, which allows 802.3 Ethernet media to carry priority information across LAN segments which are unable to signal priority (three bits long).
- The CFI field, which indicates that the Ethernet frame carries data encapsulated in native token ring format that will require translation (one bit long).
- The VID field, which uniquely identifies the VLAN on any LAN segment within the enterprise. The VID is commonly referred to as the VLAN tag (12 bits long).

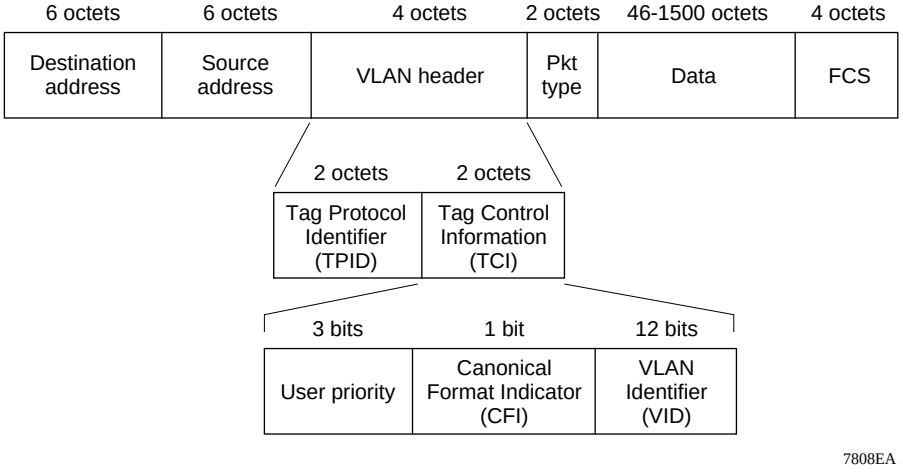

Figure 1-4. Example of 802.1q Tagged Ethernet Frame

A port cannot have two protocol-based VLANs of the same protocol type unless the port is 802.1q tagged.

In Figure 1-5, port 5 is tagged to support the IEEE 802.1q specification. When port 5 receives a packet with an 802.1q tag, the port looks for the VLAN ID contained in the tag. If the packet is meant for VLAN IP2, port 5 recognizes the VLAN ID and protocol type and forwards the packet.

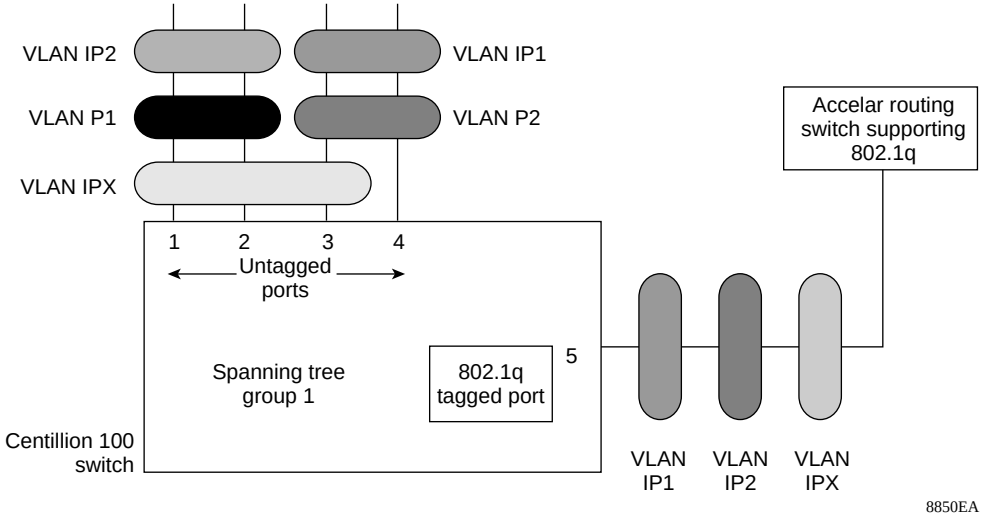

Figure 1-5. 802.1q Tagged VLANs

Conceptually, IP1 spans ports 3, 4, and 5. Moreover, port 5 can be used for multiple VLANs. Port 5 acts as a single physical link to the layer 3 Accelar routing switch for all VLANs in the spanning tree group. Without 802.1q tagging, the switch illustrated in Figure 1-5 would need five physical links to the routing switch.

Physical Ports as Members of VLANs

Physical ports are individual ports within the Centillion 50/100 or Model 5000BH switch that are associated with Ethernet, Fast Ethernet, and token ring. Physical ports within a switch may be individually joined to a VLAN within a spanning tree group. By doing so, traffic is bridged between the ports.

You can join two switches together with compatible physical ports. When joined, the VLANs associated with the physical ports in the two switches form a network in which traffic is bridged between users in the two VLANs.

Within a switch, physical ports may be joined with other physical ports to a VLAN. Ethernet and token ring ports may not be in the same VLAN.

Virtual Ports as Members of VLANs

The term virtual port describes configured ports in a VLAN that operate through any ATM physical port. A virtual port establishes connectivity with other switches and/or ATM-attached devices. An ATM virtual port is the termination port for a collection of ATM circuits—a permanent virtual circuit (PVC), permanent virtual path (PVP), or switched virtual circuit (SVC).

These circuits terminate on a Centillion 50/100 or Model 5000BH switch and operate as a single “port” in the VLAN from the perspective of bridging and spanning tree. There may be multiple virtual ports in a VLAN. There may be multiple virtual ports operating on an ATM physical port; also, there may be multiple ATM physical ports that are used by a single virtual port. A virtual port is created by the user through SpeedView.

Virtual ports of the same type may also be used to join switches.

Multiple VLANs in a Token Ring Environment

Multiple, independent source route bridge networks may be created on a given switch. Figure 1-6 shows an MCP with an internal management ring, accessible token ring VLAN 1 and VLAN 2. Traffic within the individual VLANs is bridged, but traffic between VLAN 1 and VLAN 2 is not.

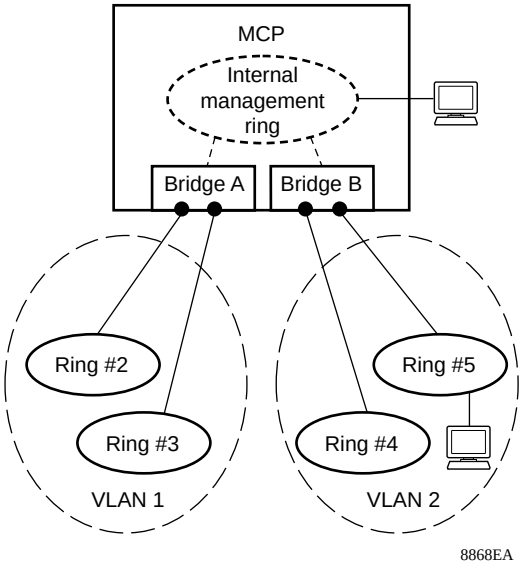

Figure 1-6. MCP with Multiple VLANs

There is still only one ring number for the MCP when multiple VLANs exist. The MCP is a member of all VLANs, and each VLAN must have a data path to the MCP.

Configuring VLANs

To have a set of physical LANs communicate within a single switch, the ports are configured to be part of a common port-based VLAN within a spanning tree group. The physical ports must be compatible with one another to be members of the same VLAN. Ethernet or Fast Ethernet ports can belong to the same VLAN, but Ethernet ports cannot be in the same VLANs with token ring ports.

For a VLAN in one switch to communicate with a VLAN in another switch, they must be connected through a physical port or virtual port. The VLANs do not have to share a name or ID for the commonly configured ports (physical or virtual) to communicate with one another.

Virtual Ring

A virtual ring is a grouping of token ring ports on a single switch that are treated as a single token ring port from the perspective of bridging and spanning tree. In particular, when a virtual ring is part of a VLAN running source route bridging (SRB), each physical port in the virtual ring is assigned the same bridge and ring number. Virtual rings are restricted to the same chassis. The rings cannot span the ATM space.

When configuring the switch for source route transparent bridging with IEEE Spanning Tree, make the following consideration: In this bridging environment, both ports in a 2-port virtual ring may be blocked by the spanning tree. Blocked ports make communication impossible between stations trying to communicate with non-source route frames between two physical rings comprising the virtual ring.

MultiLink Trunking

MultiLink Trunking (MLT) is a link aggregation technology for grouping 2 to 4 full-duplex, Fast Ethernet links to obtain higher performance between switches or between a server and a switch (Figure 1-7). MultiLink Trunking provides up to 800 Mb/s on a single aggregated trunk. The switch supports up to 32 MLT groups per switch.

The trunk acts as one logical connection between two end stations or switches. The switch distributes traffic across aggregated links providing load balancing, redundant parallel paths, and fault tolerance.

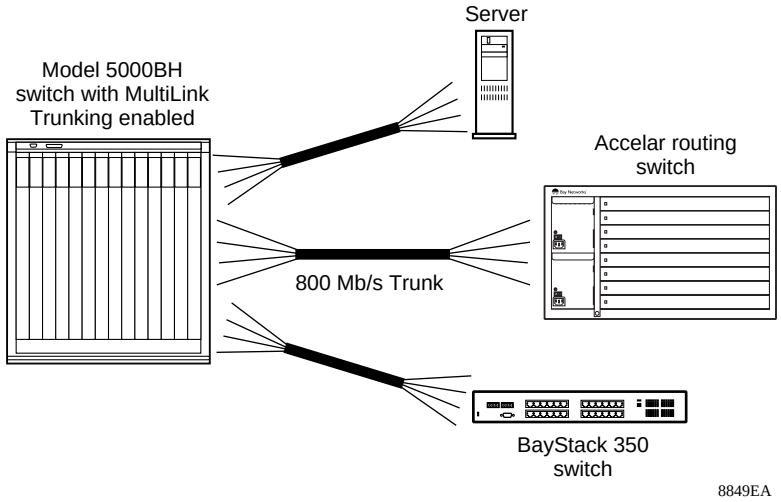

Figure 1-7. Example of MultiLink Trunking

All the ports in an MLT group are treated as one logical port. When a packet comes in with a new source address, the port with the least output traffic learns the address. If one of the ports in the MLT group goes down, all entries on the port are flushed. Other ports in the MLT group relearn the entries the next time the addresses enter one of the active ports. The switch redirects the traffic from the down link through another MLT port transparently to the user.

All the ports in an MLT group must be on the same module and have the same VLAN(s) and spanning tree domain.

A switch port can belong to one MLT group only, and when the spanning tree state or VLAN properties change on one port, the other ports must be reconfigured to match the change.

Combining MultiLink Trunking and 802.1q VLANs

Figure 1-8 shows two frame servers with MLT enabled that connect to an ATM backbone. Multiple VLANs and MultiLink Trunking are also supported between the Centillion switches and the Accelar and BayStack switches that have 802.1q enabled. The goal with 802.1q in the network is to support multiple VLANs on a single link to provide flexibility and to connect multiple VLANs to a router, switch, or server with a single physical port. Without the 802.1q tag, you would have to have a separate routing switch port dedicated for each connection to the switch. When MLT and 802.1q are combined, they provide high-performance VLAN aggregation over multiple, load-balanced links.

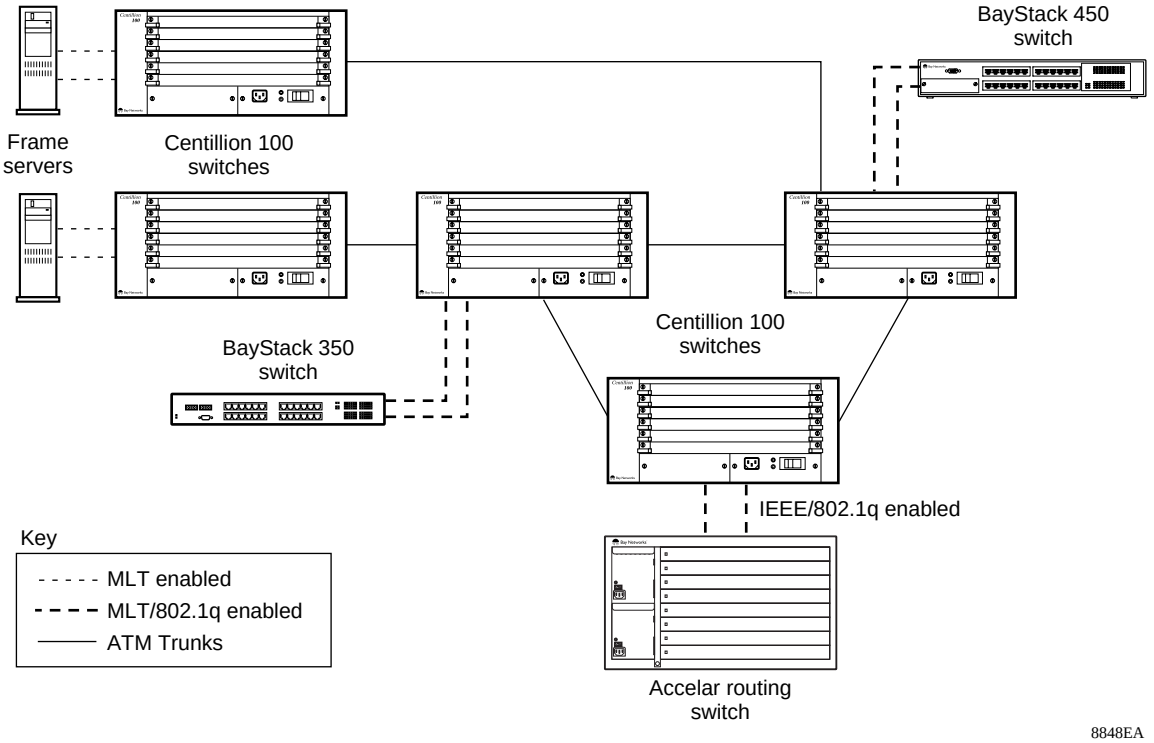

Figure 1-8. MultiLink Trunking and 802.1q in a Network

Filtering

The Centillion 50/100 and Model 5000BH switches provide powerful packet filtering. These switches can recognize any pattern within the first 255 bytes of a frame and take the action you specify for frames that match or do not match the pattern. In addition, the switches offer a NetBIOS filter template that lets you filter NetBIOS names.

You can set up filters for frames entering and leaving a physical frame port. Frames that meet the specified condition may be forwarded as usual, forwarded to additional destinations, forwarded to a monitoring port, dropped, or subjected to another filter. You can define a separate set of actions for frames that do not meet the specified condition.

SpeedView includes the following applications for filtering:

- Reducing traffic across the network
- Analyzing network performance by routing a copy of the real-time traffic to a LAN analyzer attached to a token ring or Ethernet port
- Monitoring user backbone utilization by copying traffic to a usage accounting server in addition to the intended destination
- Controlling traffic flow by redirecting frames that meet certain criteria to a specified destination instead of or in addition to the destination addresses

For more information about setting up packet filters, see Chapter 6, “Packet Filter Applications,” and refer to *Using SpeedView 4.0 for Windows*. For more information about the NetBIOS Protocol, see Chapter 7, “NetBIOS Protocol.”

Internet Group Management Protocol/IP Multicast (IGMP/IP)

IP hosts use IGMP/IP to report their multicast group memberships to immediately neighboring multicast routers. Routers send IGMP queries to all hosts, and IGMP hosts respond with IGMP reports sent to the multicast address of the group that they wish to participate in.

Centillion™ IGMP implementation treats an Ethernet port the way a router typically treats its subnets. The switch uses the information learned from IGMP activity to create a database that maps multicast groups to ports.

Note: The Centillion 50/100 and Model 5000BH switch support IGMP version 1 and 2.

Pseudo query mode (or pseudo query generation) is used when IP Multicasting is needed in the absence of a router or when IP Multicast has been administratively disabled on a router. In these instances, the switch can generate IGMP queries.

If multiple pseudo query switches exist, the switch with the smallest IP address becomes the querier. Pseudo query generation should be enabled in the absence of an IGMP query router, and it is disabled by default.

Report proxying allows the IGMP router to see a Centillion 50/100 or Model 5000BH switch as a single IGMP host with membership in multiple groups. Many IGMP hosts may connect to the switch, but report proxying allows the router to interact with only one to get the information needed. This feature prevents router performance degradation by removing the need to process a large number of IGMP reports from the end stations that are connected to the switch.

All IGMP queries are forwarded to all Ethernet ports. The switching code filters packets to the multicast Ethernet destinations after it is received from the BUS (Broadcast and Unknown Server).

In general, IGMP enables sources to send a single copy of an IP Multicast packet to the recipients that explicitly want the information. IGMP is a receiver-based concept. Receivers join a particular group and receive the traffic that is delivered to any member of that group. The sender does not maintain a list of receivers. Only one copy of a multicast message passes over any link in the network, and copies of the message are made only where paths diverge at a router.

To participate in IGMP on a local network, a host must have software that allows it to send and receive IP Multicast datagrams and to participate in multicasts that span multiple networks. The IGMP host must inform local IGMP routers of the IGMP groups that it wants to participate in. To send an IP Multicast packet, the sender specifies a destination address of a host group. In order to receive IP Multicast datagrams, the host should become a member of a Multicast group by sending an IGMP report to the router.

If more than one IP Multicast router exists in a subnet, only one router issues queries; the rest are silent. The silent routers must be advertised to the switches so that the multicast traffic will be forwarded to them. Advertising silent routers prevents the nonqueriers from being pruned.

The local switch port connected to the router is configured as a static router port so that the switch can issue the IRAP (IGMP Router Advertisement Protocol) packet to advertise the silent router to other switches. Only Centillion switches (not clients or routers) understand the IRAP packet, and the switch always treats an incoming IRAP packet as an IGMP general query. If enabled, a switch sends out IRAP packets periodically when two or more router ports are present locally.

Figure 1-9 shows a description of IP Multicast operations in a Centillion network.

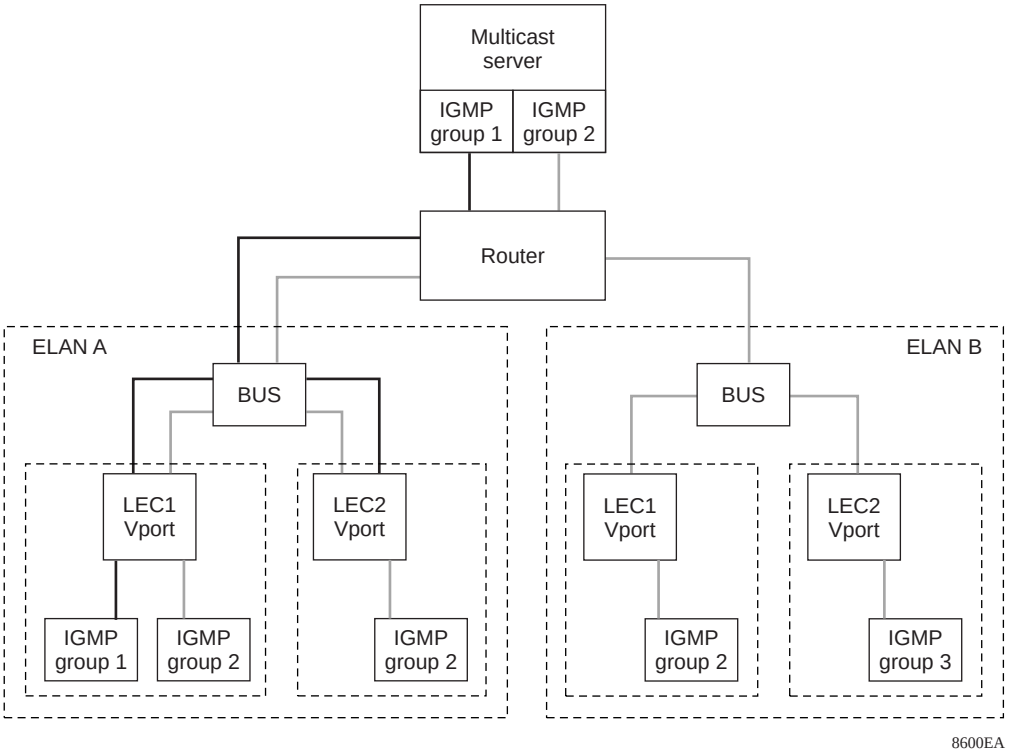

Figure 1-9. Application Behavior and Characteristics

Figure 1-9 logically depicts the behavior of IP Multicast flows from an IP Multicast server through a router and then through the BUSs of two separate ELANs to the destination Ethernet clients that have joined their respective IGMP groups.

The IP Multicast server at the top and center of the diagram is the source of the IP Multicast flows from IGMP group 1 and IGMP group 2.

In Figure 1-9, the packets flow as follows:

- A single stream of packets flows from each IGMP group within the multicast server into the router.
- The router knows that IGMP clients from each IGMP group exist in ELAN A, so the router forwards a single copy of each IGMP stream to the BUS in ELAN A.
- The BUS in ELAN A treats the multicast packets as broadcast packets and forwards them to all of the LECs registered in the ELAN.
- LEC 1 forwards the packets from the IGMP group 1 flow to all of the Ethernet clients in the switch that have joined IGMP group 1.
- LEC 1 forwards the packets from the IGMP group 2 flow to all of the Ethernet clients in the switch that have joined IGMP group 2.

Note: IGMP group 1 packets are never forwarded to ports that have not joined IGMP group 1, and the same is true for IGMP group 2 packets.

- The router knows that an IGMP client exists in ELAN B, so the router forwards a single copy of the stream from IGMP group 2 to the BUS in ELAN B.
- The BUS in ELAN B forwards the multicast packets to all of the LECs that are registered within the ELAN.
- Only LEC 1 in ELAN B forwards the packets on to any Ethernet clients.
- LEC 2 does not have a client or instance of IGMP group 2, and it does not forward any of the multicast packets from IGMP group 2 to any Ethernet clients.

LEC 2 in ELAN B has an instance of IGMP group 3. No flows exist for IGMP group 3. As a result, IGMP multicast traffic is neither received by LEC 2 nor forwarded to any Ethernet clients.

Configuration Concepts

This section describes concepts to use in your ATM configuration. This section describes the following processes:

- Using the Ethernet port on the ATM MCP (page 1-26)
- Using token ring connection modes (page 1-31)
- Signaling over nonzero virtual ports (page 1-35)
- Creating redundancy with the Centillion 50/100 and Model 5000BH switches (page 1-36)
- Connecting LANs using ATM (page 1-40)
- Using ATM LAN emulation (page 1-41)
- Using CLC mode of ATM switching (page 1-44)

Using the Ethernet Port of the ATM MCP

SpeedView supports the Ethernet port on the ATM MCP switch module. The primary IP address is the ATM MCP agent address. The secondary IP address is for the Ethernet port into the MCP agent. Packet Internet Groper (PING), SNMP, Bootstrap Protocol (BootP), TFTP, and Telnet (maximum of 4 sessions per switch) are supported through the Ethernet port on the ATMSpeed/155 MCP and Model 5724M MCP switch modules.

Only one IP gateway is supported between the primary and secondary IP addresses. If the primary IP address is using the IP gateway, then the secondary IP address must have the IP gateway set to 0.0.0.0. Also, the secondary IP subnet must be different from the primary IP subnet.

The secondary IP address can use the IP gateway if the primary IP address has 0.0.0.0 configured for the IP gateway. Also, the Ethernet port on the ATMSpeed/155 MCP and Model 5724M MCP switch modules is an MDI-X interface.

The following scenarios provide examples of how to use the primary and secondary management ports on the ATMSpeed MCP switch modules.

Scenario 1

In the configuration illustrated in Figure 1-10, the switch agent is configured to send and receive network management information on the secondary data channel. All management requests and replies take place between IP address 2 and IP address 3 on subnet B.

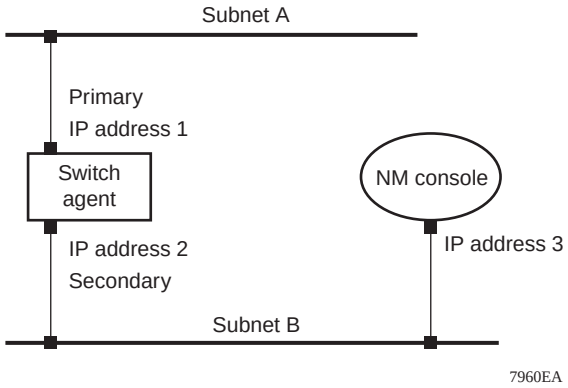

Figure 1-10. Secondary-Path Management Information Request: Out of Band

If you are watching traffic on the out-of-band Ethernet port subnet B with a network analyzer, you will see a request and the matching reply. No management traffic will appear on the switch ports subnet A.

Scenario 2

In the configuration illustrated in Figure 1-11, the switch agent is configured to send and receive network management information on the primary data channel. The default gateway is IP address 4. All requests and replies between IP address 1 and IP address 3 are routed through subnet A.

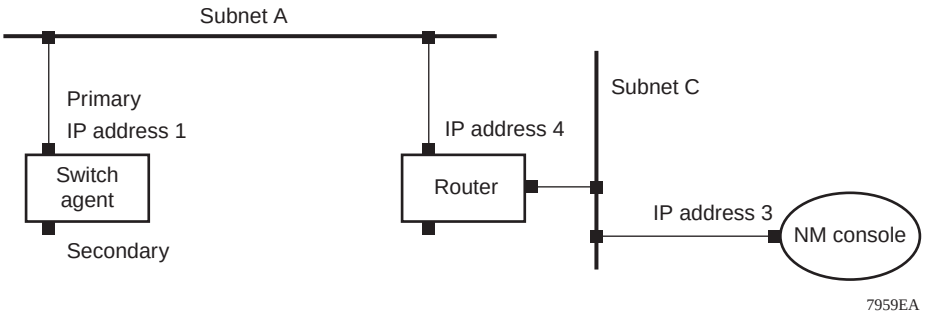

Figure 1-11. Primary-Path Management Information Request: In Band

If you are watching traffic on the switch ports subnet A with a network analyzer, you will see a request and the matching reply.

Scenario 3

Bay Networks does not recommend the configuration in scenario 3 (see Figure 1-12). This configuration results in confusing requests and responses. In this configuration, the switch agent has both the primary and secondary network management data paths configured. Management information can be requested by the network management (NM) console either out of band (secondary) or in band (primary.)

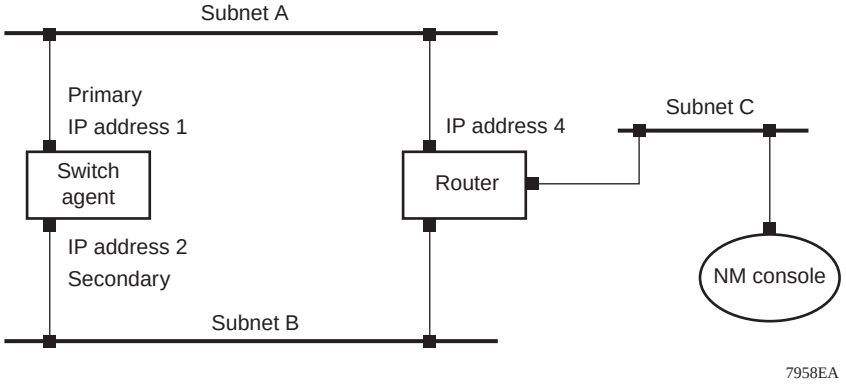

Figure 1-12. Dual-Path Management Information Request: Not Recommended

Assuming that the default gateway is IP address 4, the following situations could occur:

- a—If the network management console requests management information from the switch agent by way of IP address 1 (in band), the request is routed through subnet A. The reply from the switch agent to the network management console will be routed back along the same path using the primary IP address through the default gateway subnet A. In this case, request and reply take the same path.
- b—If the network management console requests management information from the switch agent through IP address 2 (out-of-band), the request will be routed through subnet B. The reply from the switch agent to the network management console will be routed back using the primary IP address through the default gateway subnet A. In this case, the request and reply are on different subnets.

In example b, if you are watching traffic on the out-of-band Ethernet port, subnet B, with a network analyzer, you will see a request but no reply. If you are watching on the switch port's subnet A, you will see a reply with no request. These conflicting requests and responses are very confusing.

Scenario 4

Scenario 4 (see Figure 1-13) provides a solution to the confusing configuration in scenario 3. In this network are two bridge groups (segments A and B) with a router passing traffic between them. The NM console is using the secondary management port on the Centillion 50/100 switch to obtain network management information from the switch.

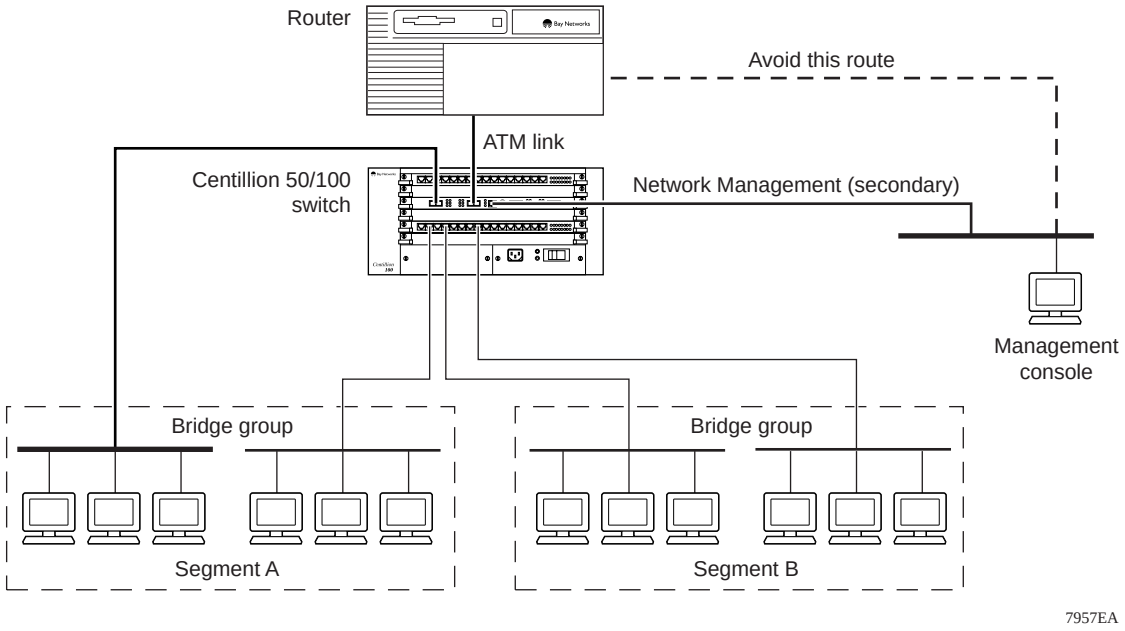

Figure 1-13. Dual-Path Management Information Request: Recommended Configuration

As explained in scenario 3, you should avoid creating a dual path between the console and the switch agent (shown as a dotted line connection in Figure 1-13). The easiest way to avoid this problem is to configure only the primary or the secondary interface. If you want to configure both to provide a redundant path for the management information, do not configure a default gateway on either the primary or secondary side. This precaution will prevent the dual-path problems described.

Using Token Ring Connection Modes

Centillion 50/100 and Model 5000BH switch modules offer the following connection modes. Each of these connection modes has several options. You can select different setting options from the drop-down list of each field within SpeedView.

- Switch Connection (Switch Connect)
- Management Wrap (Mgmt Wrap)
- Backplane Connection (Bk plane)
- Front Port Connect (Fr Port Connect)

Note: The Management Wrap, Backplane Connection, and Front Port fields apply to only the Model 5000BH switch modules.

Switch Connection Field

The Switch Connection (Switch Connect) field allows packets between the front port and the backplane to be switched. When Switch is enabled, the switch, backplane, and front port are all inserted in the same ring. This field is applicable only for ports on the Model 5000BH TokenSpeed switch modules, but it is not applicable on the Centillion 50/100 TokenSpeed modules. This field should be disabled only for troubleshooting.

Figure 1-14 illustrates how packets are switched using the Switch Connect option.

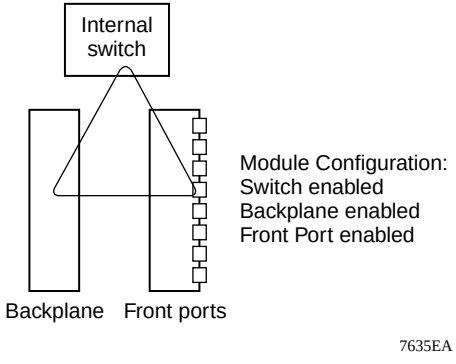

Figure 1-14. Packets Switched with the Switch Connect Option

Management Wrap Field

The Management Wrap (Mgmt Wrap) field separates the switch, front port, and backplane into three isolated rings and can be enabled for troubleshooting.

Note: When you select the Wrapped option, the change is not saved to flash memory. Once you reset the switch, the previously saved configuration is restored.

Figure 1-15 illustrates how packets are wrapped using the Mgmt Wrap option.

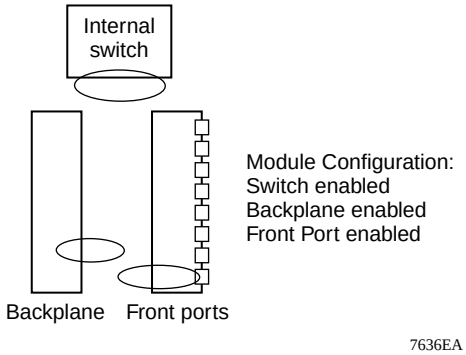

Figure 1-15. Packets Wrapped Using the Wrapped Option

Backplane Connection Field

The Backplane Connection (Bk plane) field connects or disconnects the switch to or from the backplane ring.The Bk plane field appears on the Model 5000BH switch modules only. Setting this field to enabled causes certain ports to be connected with certain backplane rings in the Model 5000BH switch.

For example, if a Model 5525 switch module (see Figure 1-16) is in slot 3 of a Model 5000BH switch with a split backplane and the backplane attachments are connected on port 5 using the Bk plane field, backplane ring 2 is connected to port 5. If the same module is in slot 9 and the backplane attachment is connected on port 5, backplane ring 7 is connected to port 5.

Note: To take full advantage of the Bk plane field, you must split the backplane using the supervisory module in the Model 5000BH switch.

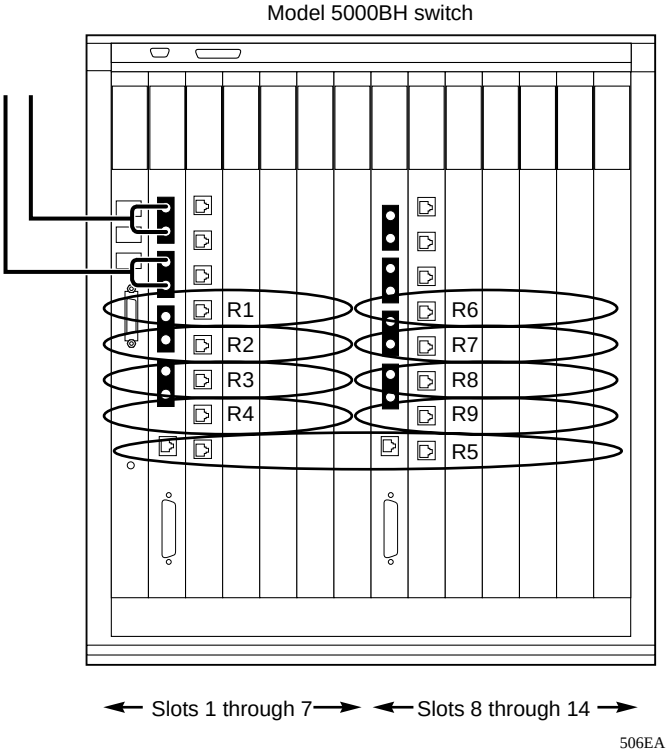

Figure 1-16. Model 5000BH Switch with a Split Backplane

Table 1-1 provides a guide to determine which slots are connected with which backplane ring.

Table 1-1. Port and Ring Assignments on a Model 5000BH Switch

Slots 2 through 7		Slots 8 through 13		
RJ-45 Port	Backplane Ring	RJ-45 Port	Backplane Ring	
			Backplane not split	Backplane split
4	1	4	1	6
5	2	5	2	7
6	3	6	3	8
7	4	7	4	9
8	5	8	5	5

Front Port Connect Field

The Front Port Connect (Fr Port Connect) field connects or disconnects the front panel RJ-45 port so that packets are either switched or not switched on a per-port basis. Figure 1-17 illustrates the different switching connections available using different variations of the Fr Port option.

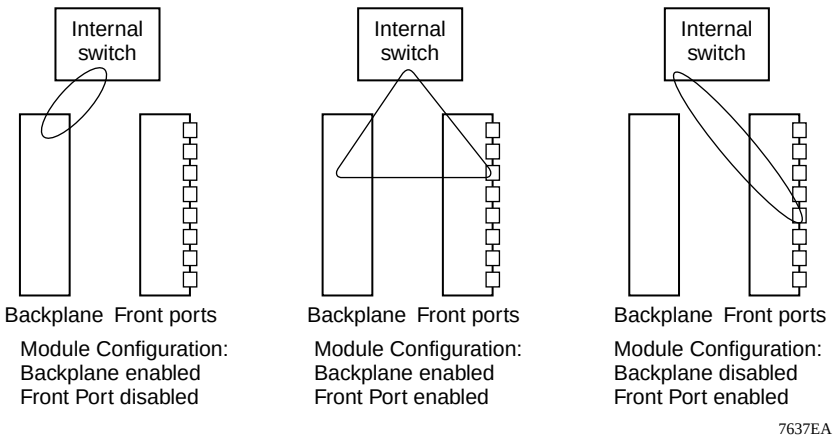

Figure 1-17. Front Port Field Options

Warning: If the front port is enabled, you must insert a cable into the front port for the backplane to attach.

Signaling over Nonzero Virtual Ports (Associated Signaling)

In WAN applications, many WAN switches use VP0 for signaling messages, which requires the ATM LAN switches in the network to signal over nonzero virtual ports. On a per port basis, the ATM LAN switch must signal on different virtual ports to establish virtual channels through the WAN to multiple LAN switches that sit behind the WAN. Centillion 50/100 and Model 5000BH switches implement multiple signaling stacks on a single physical port by allowing signaling on nonzero virtual ports as shown in Figure 1-18.

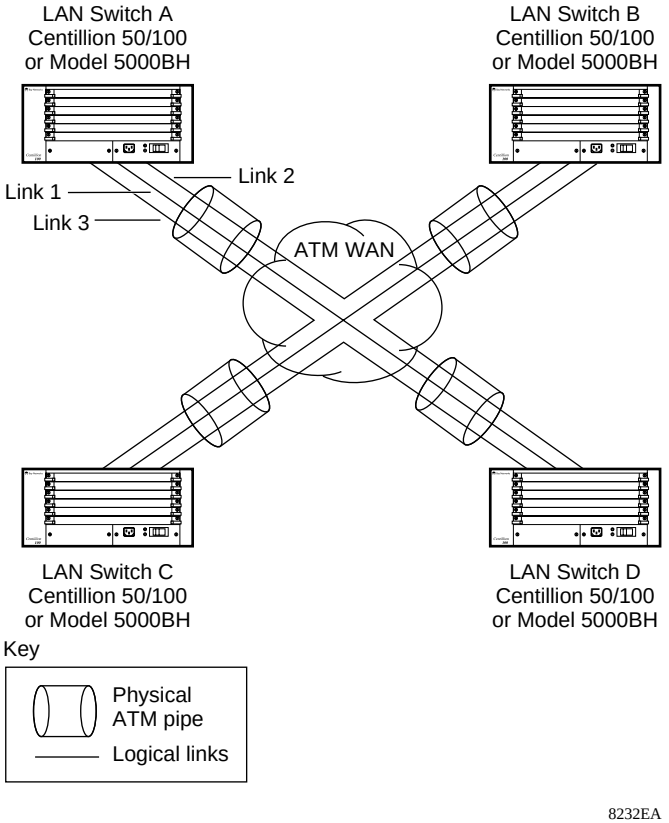

Figure 1-18. Signaling over Nonzero Virtual Ports

Creating Redundancy with the Centillion 50/100 and Model 5000BH Switches

There are several ways to accomplish redundancy within a network of Centillion 50/100 or Model 5000BH switches. The switches may be used in combination to achieve the necessary degree of redundancy. This section includes the following redundancy implementations:

- Spanning tree redundancy (this page)
- PNNI and SVC/IISP redundancy and load balancing (this page)
- LANE services redundancy (page 1-37)
- Circuit redundancy and load balancing (page 1-39)
- Redundant MCP (page 1-39)

Spanning Tree Redundancy

Redundant connections between VLANs can be accomplished with the Spanning Tree Protocol. The Spanning Tree Protocol identifies and resolves redundant paths in the network automatically. Only one path between two VLANs will be active at any one time.

PNNI and SVC/IISP Redundancy and Load Balancing

In networks using PNNI, call balancing is automatically supported.

You can create redundant connections between LEC Circuit Saver or Turbo virtual ports by using alternate-path routing for SVCs.

If an SVC between two LECs is broken, the LECs simply signal for another SVC. This call may be set up over an entirely different path as configured in the IISP routing tables for the switch. IISP defines the path over which a call request is forwarded. The IISP route entry is statically configured and may incorporate redundancy and the call balancing to a given destination.

You can use IISP to make connections to the same destination redundant and call balanced between switches. Create redundancy by making multiple IISP route entries for the same destination ATM address prefix based upon the best "length match" of the destination ATM address. If an equal match is made for two or more entries in the IISP routing table, a cost or administrative weight is used to balance calls across the matching routes. For more information about redundancy with IISP, see Chapter 5, "IISP in an ATM Network."

LANE Services Redundancy

The Centillion 50/100 and Model 5000BH switch supports LANE services redundancy, which ensures LES/BUS availability in an ELAN if the primary LES/BUS fails. You can configure two switches to provide LANE services for backup. These two switches communicate using Centillion LANE Services Redundancy Protocol (LSRP).

Every LEC belongs to a certain ELAN; however, the LEC does not know where to find LANE services until it learns the ATM address of the LES/BUS pair providing services for the ELAN.

You have two ways of providing the LEC with the appropriate LES/BUS address for the ELAN it belongs to. Either use the LECS to direct the client to the appropriate LES/BUS through LANE protocols, or manually provide the LEC with the correct LES/BUS address.

When the LEC is manually configured with a LES/BUS address, you supply an ordered list of addresses that will be tried in order in the event of failure. Up to four LES/BUS pairs are supported.

Cooperating LES/BUS Pairs

You can configure the switch to have redundant, cooperating LES/BUS pairs on an ELAN. Manually configure each LES/BUS pair with the ATM address of the other cooperating pairs for that ELAN. Each LES/BUS pair then establishes and maintains point-to-multipoint cooperating VCCs to all other cooperating pairs. For the procedure for configuring cooperating LES/BUS pairs, refer to *Using SpeedView 4.0 for Windows*.

Figure 1-19 shows an example of how LES/BUS pairs can provide redundancy and load sharing.

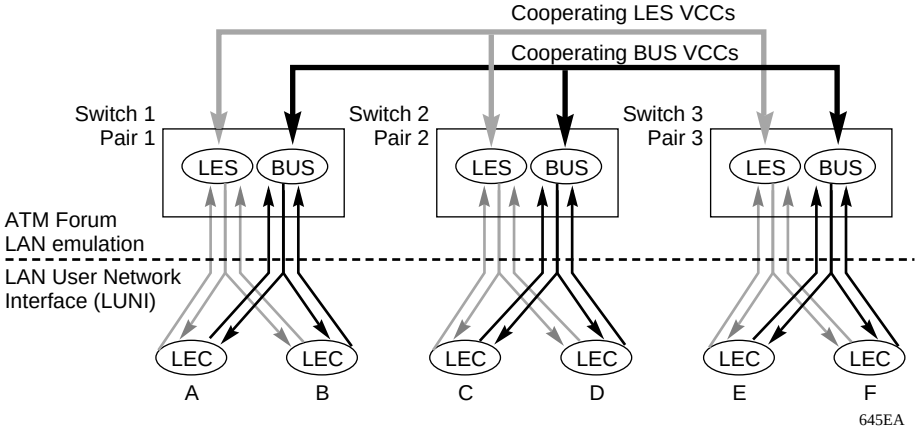

Figure 1-19. Redundant and Cooperating LES/BUS Pairs

In Figure 1-19, the LEC A and B list could be Pair 1, Pair 2, Pair 3; the LEC C and D list could be Pair 2, Pair 3, Pair 1. Each LEC establishes communication with the first available LES/BUS pair.

If you have implemented a set of redundant/cooperating LES/BUS pairs, you may specify an ordered set of LES/BUS addresses for any LEC. If the LES/BUS in the list fails, the LAN emulation client will attempt to make a connection to the next LES/BUS address in the list. The LES/BUS addresses supplied must be part of a cooperating LES/BUS group.

LECS Support for Cooperating LES/BUS Pairs

You can use the enhanced LECS support for LES/BUS pairs to make better use of the LES/BUS feature. Redundant LES/BUS pairs are configured at the LECS. When a LEC registers with the LECS, the LECS points the LEC to the next consecutive LES.

If you have configured the LEC to use a LECS, you may specify a secondary LECS address. In the event that the first LECS in the list fails, the LAN emulation client will attempt to make a connection to the next LECS address in the list.

Circuit Redundancy and Load Balancing

You can create redundant connections between CLC Circuit Saver or CLC Turbo virtual ports by defining a secondary mesh of PVCs or PVPs. For each destination end point of the primary PVC or PVP, the alternate path would ideally take a different physical connectivity path. You assign a cost, or administrative weight, to the secondary path. A lower cost signals higher priority. The lower cost path is used exclusively under normal operations. If the primary and secondary path costs are equal, traffic will be load balanced by destination MAC address across both paths. Circuit redundancy offers the quickest possible recovery while localizing the network disruption.

Redundant MCP

When enabled, the redundant MCP feature allows you to install and configure a primary and secondary (backup) ATM MCP in the same chassis. This feature is supported for ATM MCP switch modules in the Centillion 50/100 and the Model 5000BH chassis.

Note: The redundant MCP should be identical to the primary MCP. It must have the same switch software version loaded and have the same amount of RAM.

When using redundant MCP, the ATM MCPs must be in the first and second slots in the chassis; slot numbers depend upon the chassis type. For the Centillion 50/100 chassis, the MCPs must be in slots 1 and 2; in the Model 5000BH chassis, the MCPs must be in slots 2 and 3 for the left side or 8 and 9 for the right side. While in standby mode, the secondary MCP acts as a switching module with all ports available.

When both MCPs are installed, the switch uses four criteria to determine which MCP is the primary and which is the secondary. The criteria in order of importance are: priority, amount of memory, image version, and slot number.

Primary and secondary MCPs monitor each other to ensure mutual existence. If the secondary MCP detects five failed attempts to reach the primary, it will reset the switch and come up as the primary MCP, while the previous primary becomes a switching module. The failed MCP must then be swapped out to restore the backup protection.

The primary and secondary MCPs synchronize configuration changes over the control bus. If the primary MCP fails, the secondary MCP boots with the latest configuration changes received from the primary MCP.

Although MCP priority is user-configured, Bay Networks does not recommend doing so, because it may override the dynamic fault recovery. The default priority for ATM MCPs is the same. If the priority and installed DRAM and image version are the same, the MCP in the lower slot number becomes the primary MCP (slot 1 in a Centillion 50/100 switch, or slot 2 or 8 in the Model 5000BH switch).

Connecting LANs Using ATM

In a campus or a multistory building, you can interconnect multiple switches using ATM. The switch is fully capable of ATM switching, so the units can be connected directly with ATM OC-3c links; no intermediary ATM switches are required.

ATM provides the following four key advantages to your network:

- Improved throughput

  The ATM connection at an OC-3c rate provides a total capacity of 310 Mb/s per ATM link, which is three times the capacity of FDDI.

- Decreased latency

  An ATM backbone reduces queuing delays. Packets destined for a remote switch are segmented into fixed-length ATM cells. These cells efficiently interleave through the ATM backbone. Cells travel through the ATM backbone over the OC-3c ATM links with no formatting at each hop. Large packets never block small packets, and transmissions are not slowed by low-speed links. Logically, an ATM network is a one-hop network.

- Compatibility

  The ATMSpeed/155 module is ATM Forum compliant, which means it can communicate with other ATM Forum compliant switches or carrier services.

- Fiber savings

  Distributed token ring backbones often require one fiber pair per bridge or router. In contrast, when token rings are consolidated into a single ATM backbone, only one fiber pair is needed between any two switches.

Using ATM LAN Emulation

The switch can be configured to implement LAN emulation as defined by the ATM Forum LANE specification. When a virtual port is configured to act as a LAN emulation client, either the Circuit Saver or Turbo method of switching can be selected.

LAN emulation provides the following features:

- Extended shared-media domains are supported.
- Emulated LANs are easily configured through network management.
- Communication through the switch to user communities is transparent.
- The conversion of LAN packets in ATM cell format is supported.
- The physical network infrastructure does not need to be changed; the prediction of logical network structure is eliminated.
- One switch can support multiple clients.

The switch can be configured as a fully functional, standards-based ATM switch, supporting UNI 3.0 and UNI 3.1 on the network side, PNNI, and IISP. ATM devices such as servers, routers, and other ATM clients can register with the switch and signal for connections to other ATM devices.

A simple configuration is a centralized star topology, as shown in Figure 1-20.

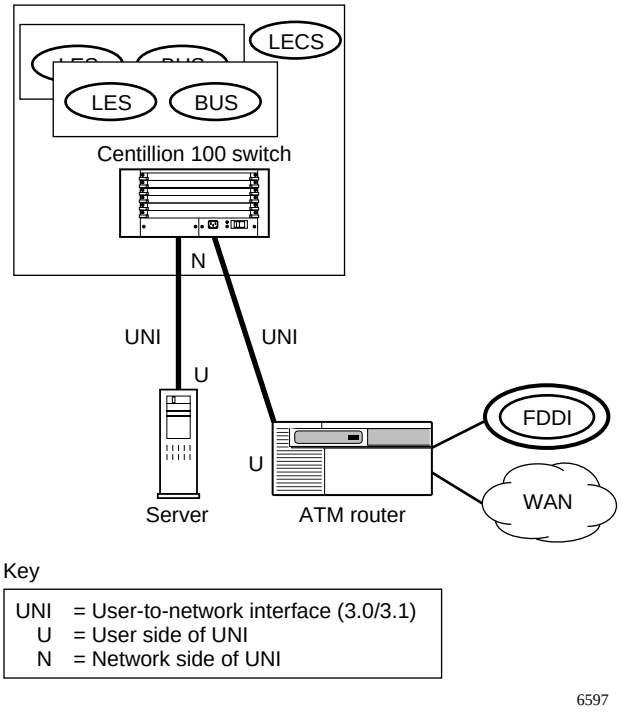

Figure 1-20. The Switch in a Centralized ATM Switch Configuration

In a centralized star topology, the edge switches can stay as ATM clients and the "network" is really only the central switch.

Another approach is to expand the “network” to the periphery and have multiple switches be the network. Switch-to-switch communications are performed using IISP and/or PNNI. An example is shown in Figure 1-21.

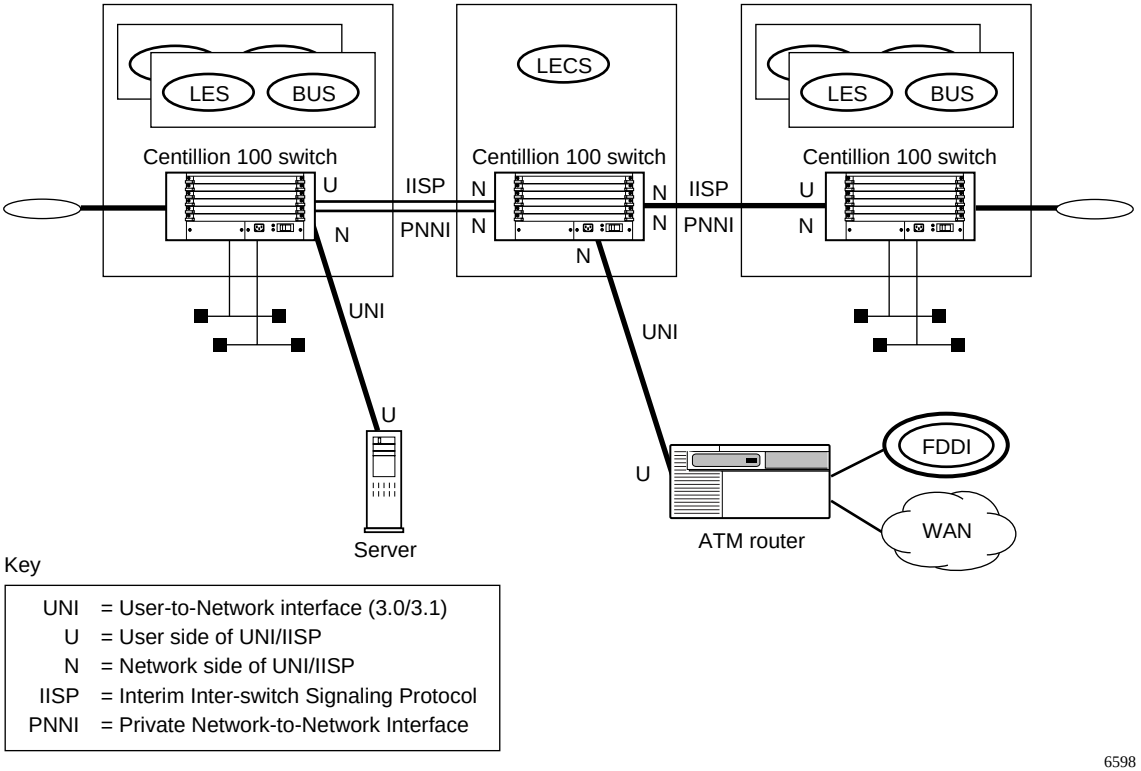

Figure 1-21. The Switch in a Distributed ATM Switch Configuration

IISP is a static routing protocol, and PNNI is a dynamic link routing protocol; either protocol can be used to allow two peer ATM switches to signal an SVC end to end, across multiple switches. Using IISP, SVCs can be created through many switches, allowing ATM native devices anywhere in the network, not just on the center switch.

On the LANE services side, LECS and LES/BUS pairs can exist on any switch in the network. Also, they can be distributed among multiple switches to provide redundancy and call balancing.

Using CLC Mode of ATM Switching

This section describes how virtual ports can be configured into logical LANs in the ATM network to suit particular performance and connectivity objectives.

As described earlier, a Centillion LAN client (CLC) can be configured for either the Turbo or Circuit Saver forwarding method. When CLCs using Turbo mode are connected (using PVPs), they automatically form a GIGArray LAN to the ATM network. When CLCs using Circuit Saver mode are connected (using PVCs), they automatically form a CLC Circuit Saver LAN.

Table 1-2 summarizes the differences between the two types of switching.

Table 1-2. CLC Circuit Saver versus Turbo

Features	Circuit Saver	Turbo
Performance	High	Very high
Circuit connections	Virtual circuits	Virtual paths
Redundancy	Supported	Supported
Load sharing	Supported	Supported

The following sections provide examples of Circuit Saver and Turbo switching.

Using CLC Circuit Saver Switching

Direct Circuit Saver switching over ATM between two switches is provided when those switches are configured in the same Circuit Saver LAN. Configuring a single switch to participate in more than one Circuit Saver LAN provides flexibility in the trade-off between performance (all switches on the same Circuit Saver LAN—requires more circuits) and configuration ease (switches configured in multiple Circuit Saver networks—requires fewer circuits). The following two figures illustrate this trade-off.

Figure 1-22 illustrates the case where the user configures all switches in the same Circuit Saver LAN.

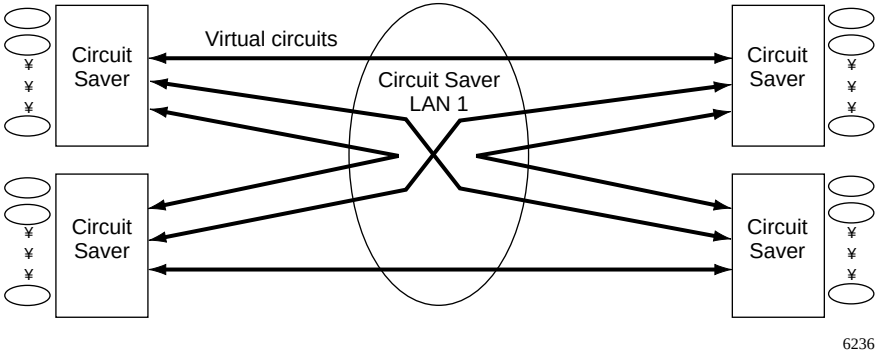

Figure 1-22. Single Circuit Saver LAN (Token Ring Example)

This configuration allows direct ATM forwarding between any two of the switches with better performance, but it requires more circuits to be set up.

Figure 1-23 illustrates an Ethernet configuration where the user wants the fewest circuits to be used. With fewer circuits, performance is lower.

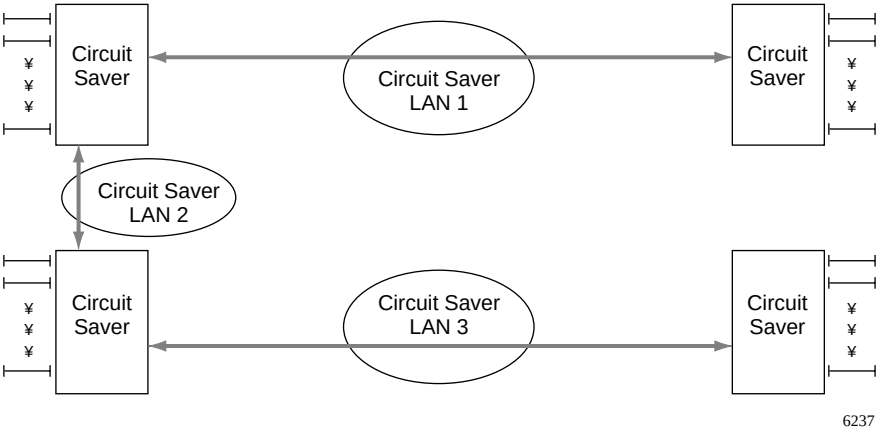

Figure 1-23. Multiple Circuit Saver LANs (Ethernet Example)

This configuration still allows the entire environment to be bridged; however, some forwarding is done via intermediate switches along the path. In Figure 1-23, several of the switches are configured to belong to more than one Circuit Saver LAN.

Up to 16 Circuit Saver virtual ports (sets of ATM circuits) can be configured on a single switch. The switch can bridge between a GIGArray LAN and multiple Circuit Saver LANs.

Using Centillion GIGArray Switching

A switch becomes part of a GIGArray when the user creates a CLC on that switch and then configures a PVP to each of the other CLCs in the GIGArray. When configuring the client as the end point of one of these PVPs, the user is required to enter a nonzero virtual port identifier (VPI) value (between 1 and 255). In addition, the user can optionally enter a virtual circuit identifier (VCI) value, which specifies the starting value of the automatically assigned VCIs within that PVP.

To configure a set of CLCs (on different switches) to be part of the same GIGArray (connects to the same GIGArray LAN), a full mesh of PVPs must be set up among them. Between any two switches in the same GIGArray, at least one point-to-point PVP (with a nonzero VPI) must be set up. Multiple PVPs can be configured between any two switches for redundancy and load sharing purposes. Virtual circuits within the PVP are set up automatically between switches in the same GIGArray.

Note: The requirement for a “full mesh” of connections within a GIGArray does not imply that the participating switches need to be connected together in a full mesh of physical links. The full mesh requirement refers only to PVPs, which can be switched transparently through other switches between end points.

Redundancy and Load Sharing

Between any two switches in the same Circuit Saver or GIGArray LAN, two or more PVCs or PVPs, respectively, can be configured to provide redundancy. When more than one PVP exists (on the same or different physical ATM ports) as part of the same GIGArray virtual port going to the same virtual port on the other end, the switch uses a PVP with the lowest cost as the primary link. Any PVP configured with a higher cost is treated as a backup link. If a lowest-cost PVP fails, then the next set of lowest-cost PVPs will be used. Users see little or no disruption of service. A load-splitting algorithm is employed when redundant PVPs have the same user-configured cost parameter. In a high-traffic situation when one physical ATM port is saturated, additional PVPs are needed to divide traffic between multiple physical ATM ports.

Load sharing is always enabled when multiple PVPs between two switches have the same user-configured cost. Load sharing can be disabled or bypassed by configuring the PVPs between two switches with different costs. The choice of path is done on a per-station connection basis; no load sharing is done on a packet or cell basis.

Full Mesh Connectivity

One useful configuration technique is to use Circuit Saver LANs to interconnect multiple GIGArrays. By interconnecting multiple GIGArrays, the user can optimize performance among the switches with the highest interswitch communication needs while conserving circuit resources (for greater scalability) for communication between GIGArrays.

When using CLC Circuit Saver and CLC Turbo LANs, each individual member of the CLC Circuit Saver or CLC Turbo LAN group must have circuits defined with every other member. All Centillion 50/100 and Model 5000BH switch virtual ports in a CLC Circuit Saver or CLC Turbo LAN should have a complete mesh of configured PVCs or PVPs, respectively.

In the example shown in Figure 1-24, four GIGArrays are created.

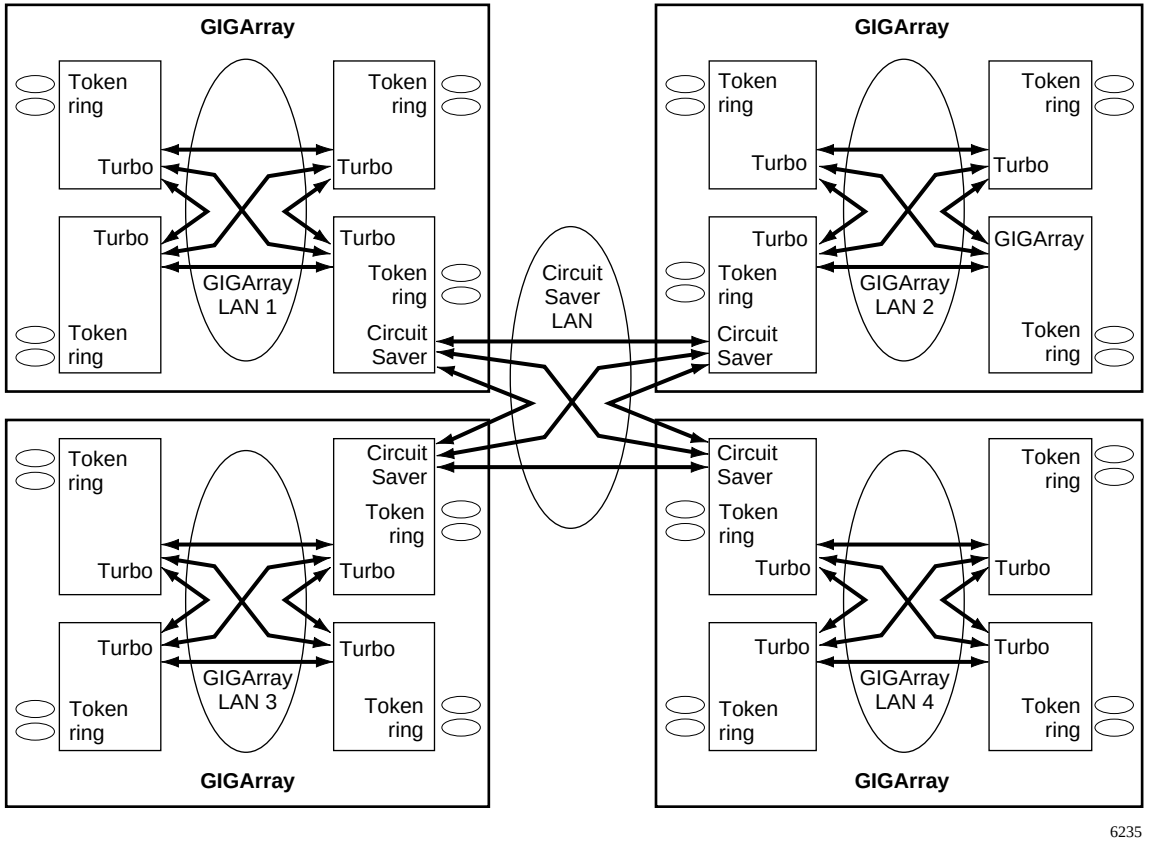

Figure 1-24. GIGArrays Connected Using Circuit Saver Bridging

The GIGArrays are interconnected by one Circuit Saver LAN. If all 16 systems were connected using GIGArrays, 120 virtual paths would be required to achieve full-mesh connectivity:

[(16 * 15) / 2] GIGArray PVPs = 120

By separating the systems into four GIGArrays and using Circuit Saver mode to interconnect the GIGArrays, only 24 virtual paths and 6 virtual circuits need to be configured to achieve full-mesh connectivity:

{[(4 * 3) / 2] GIGArray PVPs * 4 GIGArrays} + [(4 * 3) / 2] Circuit Saver PVCs = 30

Implementing LANE Services

You should have a fundamental understanding of the LANE standard and how a LEC interacts with the LANE services (LECS, LES, and BUS) in order to understand the implementation choices and trade-offs for LANE services on the switch. Chapters 1 through 3 of the ATM LANE specification provide an excellent overview of these concepts. The specification may be found at the ATM Forum's Web location.

The choices for implementing LANE services affect how you configure the services and all the LECs participating in the network, CLC and LANE virtual ports, and other ATM-attached devices such as VNRs and servers.

You have the following two choices:

- Use a LAN emulation configuration server (LECS) and configure each LEC in the network to do so.
- Bypass the LECS by providing the specific LES/BUS address for the chosen emulated LAN when configuring the LEC.

Using the LECS is easier; but if a failure occurs, it takes longer for the network to reconfigure.

Using a LECS

Using a LECS as a first point of contact for the LAN emulation clients has some advantages. Because the basic function of the LECS is to provide the LEC with the appropriate ATM address of the LES/BUS for an ELAN, the LES/BUS ATM address only needs to be configured in the LECS Configuration window. If the LES/BUS address for an ELAN changes as a result of reconfiguration, only the LECS must be reconfigured.

Using LECS Redundancy

If you choose to provide multiple LECS ATM addresses that the LEC may use, make sure that all configurations for the respective LECSs are equal. The ELAN name and LES ATM address must be the same in both LECSs. Different name and address information implies that clients talking to a different LECS are placed on a different ELAN. To take advantage of this redundancy, configure the LEC to use more than a single LECS. LANE virtual ports allow this feature.

Note: ELAN names are case sensitive. ELAN1 is different from elan1.

Using the LES/BUS Directly

To direct a LEC to use a LES/BUS directly, configure the LEC to use the ATM address of the specific LES/BUS desired. One advantage of making the LES/BUS configuration directly is that the LECS configuration step is saved during startup and failure-mode recovery.

Using LES/BUS Redundancy

The cooperating LES/BUS feature within the Centillion 50/100 and Model 5000BH switches allows you distributed and redundant LES/BUS functionality for a given ELAN. To achieve redundancy, configure LES/BUS entities within separate switches. Within a switch, identify the other LES/BUS pairs in other Centillion 50/100 and Model 5000BH switches and assign a unique server ID to the LES/BUS pair cooperating with them.

To take advantage of redundancy, configure the LEC to use an ordered list of LES/BUS addresses in the event of failure. LANE virtual ports allow this feature. To take advantage of the distributed capabilities, point different LECs to different LES/BUS pairs to distribute the load.

Recommendations and Tips

Combining the LEC function and LANE services on a switch reduces LE-ARP traffic because the information is available to the switch's LEC without communicating over an ATM link.

All of the NSAP addresses for LEC, LES, BUS, and LECS must be unique. When these components are within the same switch, uniqueness can be ensured by making your changes in the End System Indicator (ESI), selector byte portion, or both of the NSAP address, which are the last 7 bytes of the 20-byte NSAP. SpeedView automatically creates unique LEC, LES, BUS, and LECS addresses.

If you are using the LECS function, configure the same ELAN name within both the LECS and respective LES/BUS configuration windows. The ELAN name is transmitted during both the configuration and the registration process and for compatibility.

If you are supplying multiple LES/BUS addresses to a LANE virtual port, make sure the ATM addresses you have provided are members of a defined, cooperating LES/BUS group. Not doing so may lead to incorrect configuration or failure of a LANE virtual port to establish connectivity. Do not identify more than one cooperating LES/BUS pair for a particular ELAN on the same Centillion 50/100 or Model 5000BH switch.

If several switches have LANE virtual ports configured, it is best to configure only a subset of these for LANE services. For example, even when using redundant LES/BUSs, usually it is not necessary to configure more than two or three switches for LANE services.

Remember to configure IISP routing tables or enable PNNI for all the switches that have LANE virtual ports and/or LANE services configured. You must do this so that switched virtual circuits may be signaled between the switches.

The standard image supports LANE and IISP and requires a minimum of 10 MB of memory. The advanced image supports LANE, PNNI, and MPOA and requires a minimum of 16 MB of memory. An ATM MCP with 26 MB of memory can accommodate up to 32 LES/BUS instances. The maximum number of ATM ports that can be configured with 26 MB of memory is 24 ports. LAN Network Manager (LNM) configuration also requires a minimum of 26 MB of memory on the ATM MCP.

Table 1-3 shows the memory recommendations for the ATM MCP.

Table 1-3. ATM MCP Memory Requirements

Configuration	Number	Memory
Standard Image		10 MB Minimum
Advanced Image		16 MB Minimum
LES/BUS	Up to 32	26 MB
ATM ports	Up to 24	26 MB
LNM		26 MB

When you are defining LECs for a Centillion 50/100 or Model 5000BH switch, no two LECs should be joined to the same ATM Forum emulated LAN, whether they are on the same VLAN or on different VLANs. Only one internal LEC should be configured per ATM Forum emulated LAN in a Centillion 50/100 or Model 5000BH switch.

LAN Network Manager and Ring Parameter Server

The Centillion 50/100 and Model 5000BH switches are manageable from IBM LAN Network Manager (LNM). This tool is used at many large token ring sites and ensures compatibility with existing IBM bridges and token ring management architectures. The Centillion 50/100 and Model 5000BH switches function like multiple IBM bridges on the network. This feature enables LNM to access each ring to which the switch is connected and to view the switched network topology as if an IBM bridge were in use.

Every bridge port on the switch is connected to a token ring. A bridge on each port connects between the external rings and the internal switch fabric.

Conceptually, packets move from a station on an external ring through the bridge port to the internal switching fabric (which LNM displays as a ring) and out another bridge port to the destination station on another external ring (Figure 1-25).

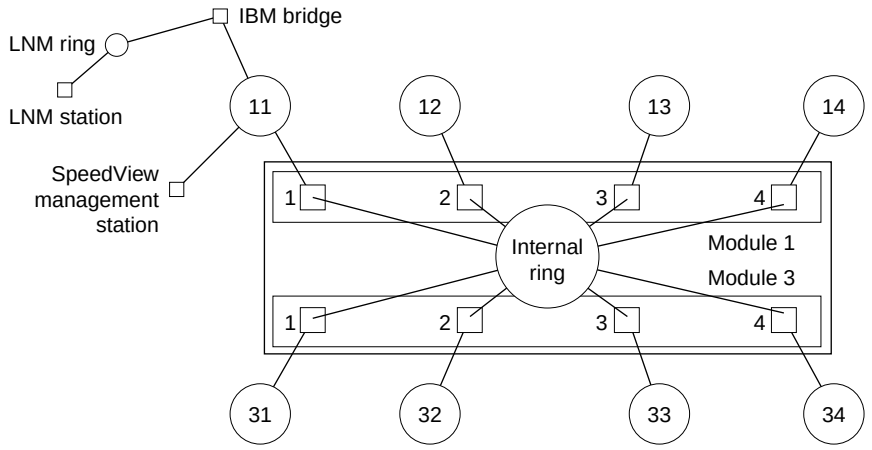

Figure 1-25. Internal and External Rings Connected by Bridge Ports on a Switch

Each port has two addresses. One is the burnt-in MAC address assigned at the factory. The other is the locally administered version of the burnt-in MAC address.

For example:

The address 000500082940 is the burnt-in MAC address.

The address 400500082940 is the locally administered version (40 in the first byte indicates LAA format).

Both addresses are available for use by the port. By default, on a token ring MCP, the MAC address of the switch (displayed in the Admin MAC field in the System window of SpeedView) is in locally administered address (LAA) format.

With LNM, you can display the internal ring as one large ring or display it as multiple rings corresponding to bridge groups defined with SpeedView (Figure 1-26 and Figure 1-27).

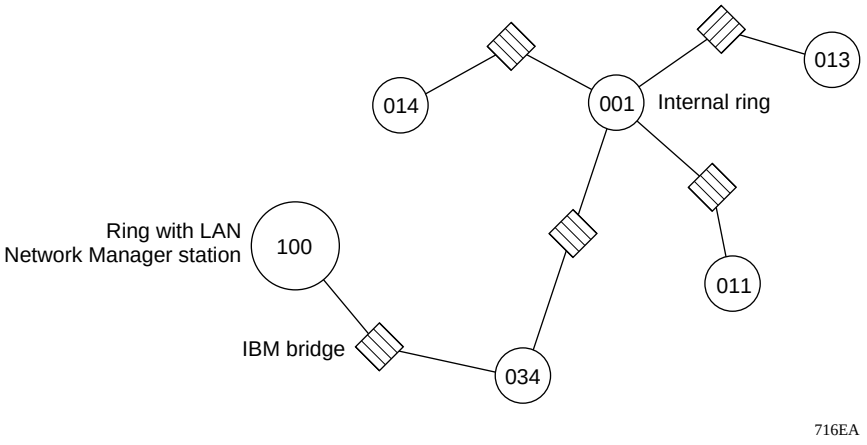

Figure 1-26. All External Rings Connected to One Internal Ring

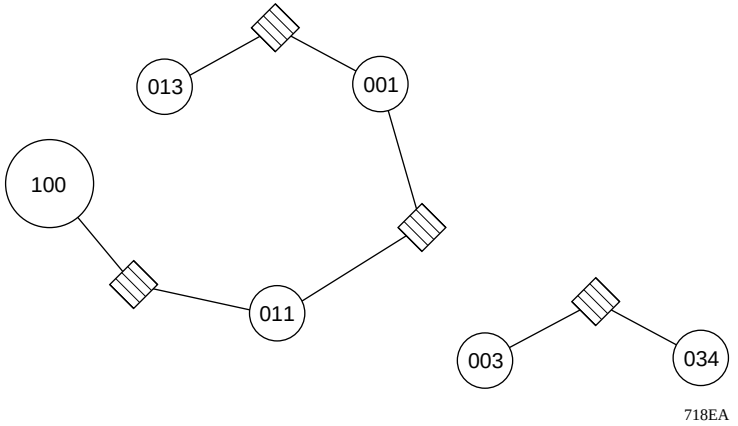

Figure 1-27. Each Bridge Group Displayed as a Separate Internal Ring

Using LNM to Display the Switch

To have LAN Network Manager display the Centillion 50/100 or Model 5000BH switch, you need to know the locally administered address (LAA) of the switch. The LAA appears in the System window of SpeedView. If you know the main address of the switch, you can add a bridge definition to LNM. Each port LAA is displayed through SpeedView in the statistics windows.

Note: The LAA can also be viewed with the CLI *show port* command.

After you connect to this bridge, you can click on the internal ring to learn the addresses of other switch ports that are inserted. You can define more bridges using these MAC addresses.

You can find the address of a bridge port from another bridge port on the same module because the burnt-in MAC addresses are assigned in sequential order. If the address of module 1, port 1 is 000500812520, then the address of module 1, port 3 is 000500812522. Each switch module has a different range of MAC addresses.

Managing Virtual Ports Using LNM

You can also manage virtual ports using LNM. You must add a “Bridge Definition” to the LNM database for each virtual port. In addition, you need a unique MAC address to add the bridge definition of each virtual port. To obtain a unique MAC address, use the Admin MAC Address displayed in the System window and create a unique MAC address for the bridge definition for the virtual port.

Note: You must have an ATM MCP switch module to manage your network across ATM links using LNM.

To create a unique Admin MAC Address:

1. **Find the Admin MAC Address displayed in the System window.**

   For example, a generic Admin MAC Address is 4005xxYYYYYY.

2. **Replace "xx" with the Vport ID of the virtual port that you want to manage using LNM.**

   For example, if the Admin MAC Address is 40050008ABCD and if the VPort ID is 1, the MAC address to enter in the LNM "Add Bridge Definition" field is 40050108ABCD.

In transparent bridging mode, the ATM cloud is displayed as ring F00 (hex). In SRB/SRT mode, the ATM cloud is displayed as a ring you assigned to the virtual port using SpeedView.

Configuring the Ring Number

When you manage your network through token ring, the ring number carrying network management traffic for each Centillion 50/100 or Model 5000BH switch in a source route bridge (SRB) or source route transparent (SRT) bridged network must be unique.

The management ring number assigned to each switch must not be used anywhere else in the entire token ring network.

Ring numbers 0 and FFF are reserved and cannot be used. The range supported for the internal ring number is 1–FFE (hex). If you are using LNM to manage your token ring network, use a ring number in the range of 1–FDE (hex).

You can use FFF (hex) as an external ring number.

Ring Parameter Server

The Ring Parameter Server (RPS) provides initialization information to new stations and forwards registration information to LAN Network Manager. For more information about RPS, refer to *Token Ring Architecture Reference* by IBM Corporation.

The switch will function as the RPS when LNM is enabled. The RPS provides the configured ring number on a port to a newly attached station. To disable RPS, turn off LNM.

Chapter 2
Configuration Examples

This chapter presents examples of how the Centillion 50/100 and Model 5000BH switches can be used in different network configurations to improve network performance. These examples provide a foundation for developing more complete networking solutions.

This chapter provides the following sections:

- Token ring solutions (this page)
- Ethernet solutions (page 2-13)

Token Ring Solutions

This section provides solutions for the following token ring network needs:

- Giving a server a dedicated port (page 2-2)
- Segmenting a token ring LAN (page 2-2)
- Replacing bridges or routers (page 2-4)
- Concentrating LANs into a router (page 2-5)
- Reducing the number of source route hops (page 2-8)
- Collapsing a token ring backbone (page 2-9)
- Using distributed backbone switching without LANE (page 2-10)
- Using distributed backbone switching with LANE (page 2-12)

Giving a Server a Dedicated Port

Giving a server a dedicated port on the switch improves server response and increases throughput capacity. When the server resides on a token ring, the switch provides the necessary token ring concentrator functions to enable you to connect directly to the switch with a lobe cable.

Giving a server a dedicated port on the switch need not increase the source route hop count. You can maintain or even decrease the source route hop count by configuring the server on the same virtual ring as its communicating stations.

Segmenting a Token Ring LAN

Excessive traffic on a token ring LAN slows response time to all stations. Using the switch, you can segment a congested LAN to create multiple LANs, thereby increasing the average bandwidth available to each device.

Token ring LANs may be segmented and connected to the switch using either the Ring In (RI) or Ring Out (RO) port or a concentrator port on a token ring concentrator.

Figure 2-1 and Figure 2-2 show a token ring LAN before and after segmentation.

Figure 2-1 shows a token ring LAN segmented using a concentrator port.

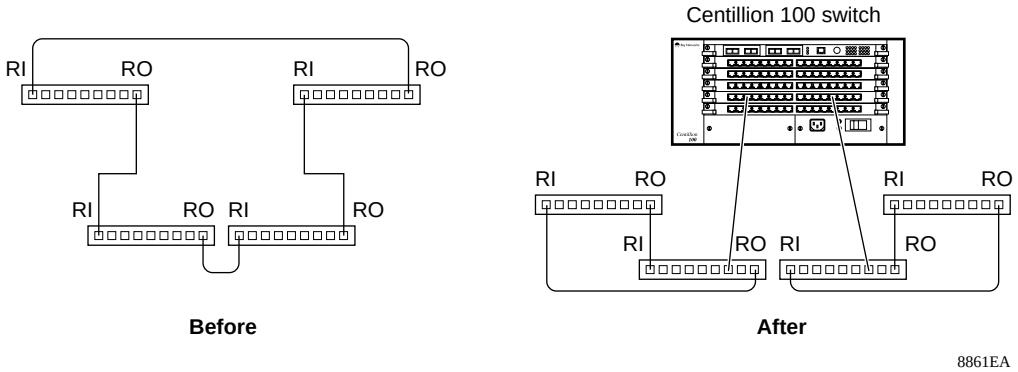

Figure 2-1. Segmenting a Token Ring Using a Concentrator Port

Figure 2-2 shows a token ring LAN segmented using the RI or RO port. The switch is preconfigured for microsegmentation.

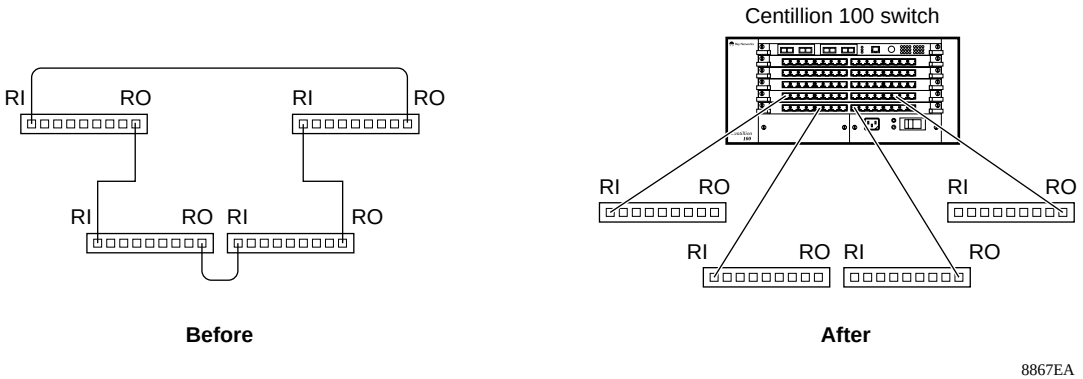

Figure 2-2. Segmenting a Token Ring Using the RI or RO Port

Replacing Bridges or Routers

The switch offers two benefits as a replacement for bridges and routers. First, it provides additional ports at lower equipment cost (Figure 2-3). Second, it increases the aggregate bandwidth of the LAN backbone. Figure 2-3 shows a Centillion 50/100 switch replacing source route bridges. This example also applies to the Model 5000BH switch.

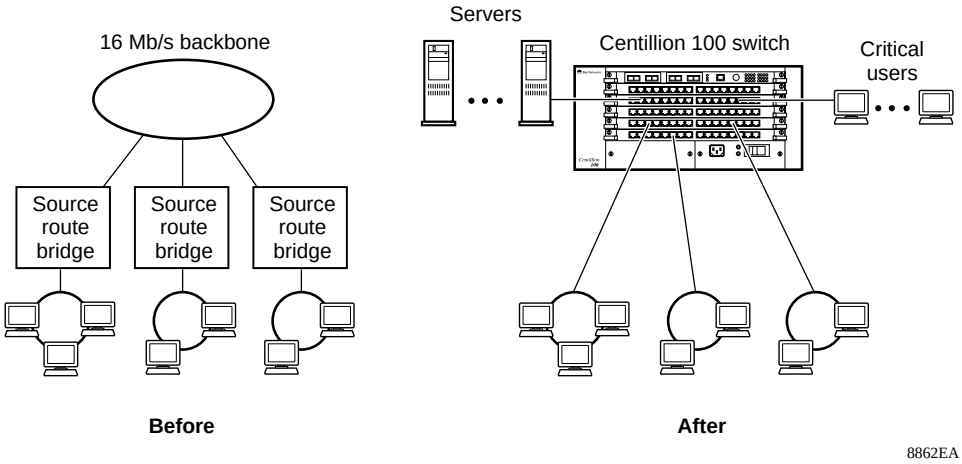

Figure 2-3. Replacing Source Route Bridges or Routers

The switch is preconfigured as a multiport transparent or source route bridge for replacement of multiple two-port bridges.

Concentrating LANs into a Router

The switch can relieve an overloaded router of port and processing burdens. The router in the network illustrated in Figure 2-4 has no capacity for additional rings, which prohibits the network manager from segmenting existing rings or adding new rings.

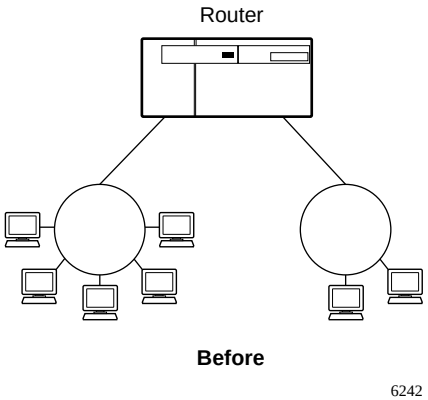

Figure 2-4. Token Ring Network with No Growth Capacity

Figure 2-5 and Figure 2-6 illustrate two approaches to using the switch to add rings and free router ports.

By reconfiguring the network as in Figure 2-5, the source route bridging (SRB) processing may be completely offloaded from the router to the switch. Offloading the SRB processing frees valuable memory and processing cycles on the router, providing better performance and longer equipment life. This example also applies to the Model 5000BH switch.

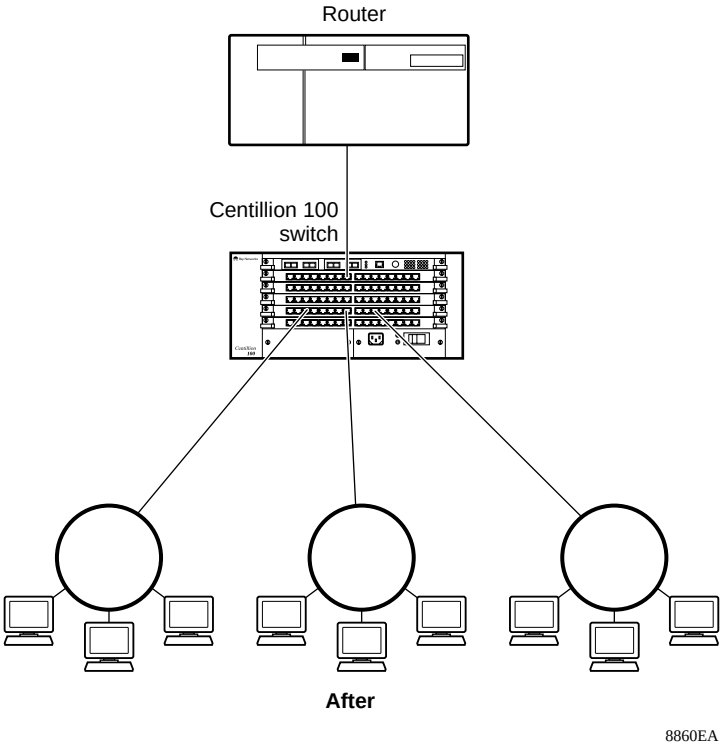

Figure 2-5. Switch Concentrating Multiple Rings as a Single Bridged Subnet

When the switch concentrates LANs into a router, its virtual ring and Ethernet virtual LAN partitioning capability lets you define one or more router subnets. For more information about configuring virtual rings, see Figure 2-6 and refer to *Using SpeedView 4.0 for Windows*. Also, by placing ports into the same VLAN, you can similarly create token ring virtual LANs. The example in Figure 2-6 also applies to the Centillion 50/100 switch.

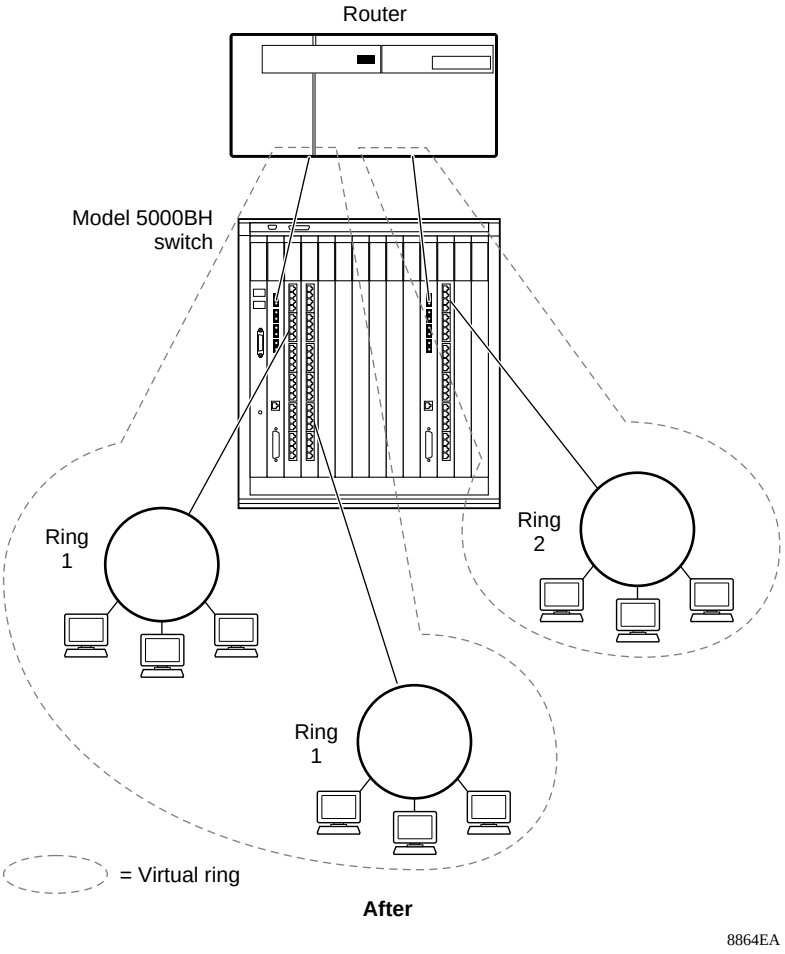

Figure 2-6. Switch Concentrating Multiple Rings as Multiple Routed Subnets

Reducing the Number of Source Route Hops

Source routing imposes a hop-count limit on token ring topologies. The limit for most installed equipment is seven hops; for certain new adapters, bridges, and applications, it is 13 hops.

Note: For 13 hops to be the limit, all devices in the route must support 13 as the hop-count limit. This limit may inhibit your ability to segment congested rings to improve performance.

The switch reduces the number of source route hops by allowing you to group multiple physical segments into a single virtual ring.

Figure 2-7 shows an example of grouping segments into a single virtual ring.

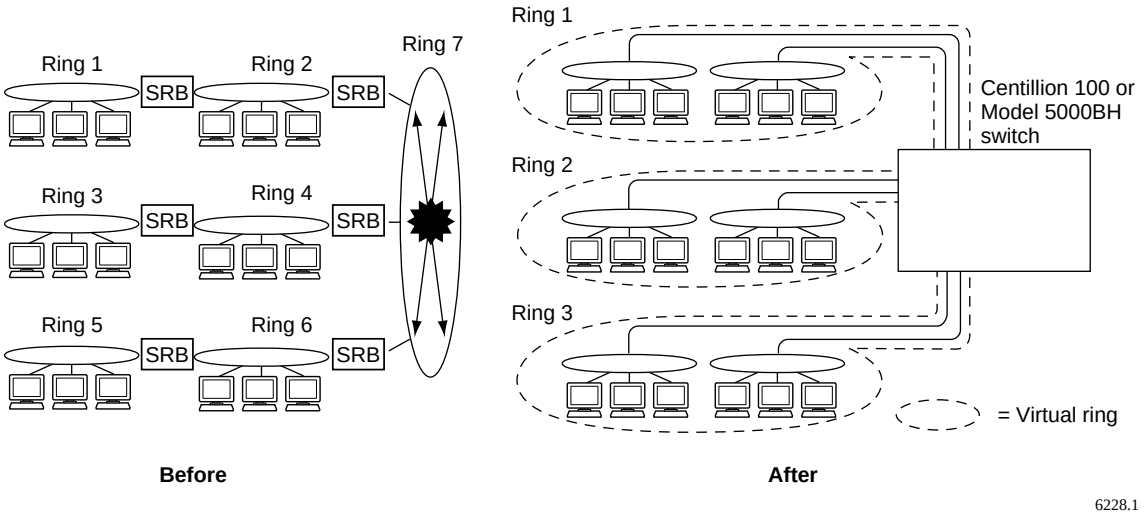

Figure 2-7. Reducing the Number of Source Route Hops

Each virtual ring in a source routed network appears as one ring with a single ring number. Each virtual ring appears as a single subnet to a router.

With fewer hops, network performance improves for two reasons:

- You can more freely segment congested rings.
- You decrease latency with each hop you eliminate.

Collapsing a Token Ring Backbone

Backbone rings can become bottlenecks as traffic between LAN segments increases. The switch can serve as a collapsed backbone, increasing the aggregate bandwidth of the backbone to multiple gigabits per second.

Figure 2-8 shows an all token ring network that collapsed many ring segments from various hubs to the network center while also providing “big pipes” to servers. A Centillion 50/100 switch could be used in the network center as well to provide the same backbone and switching solutions.

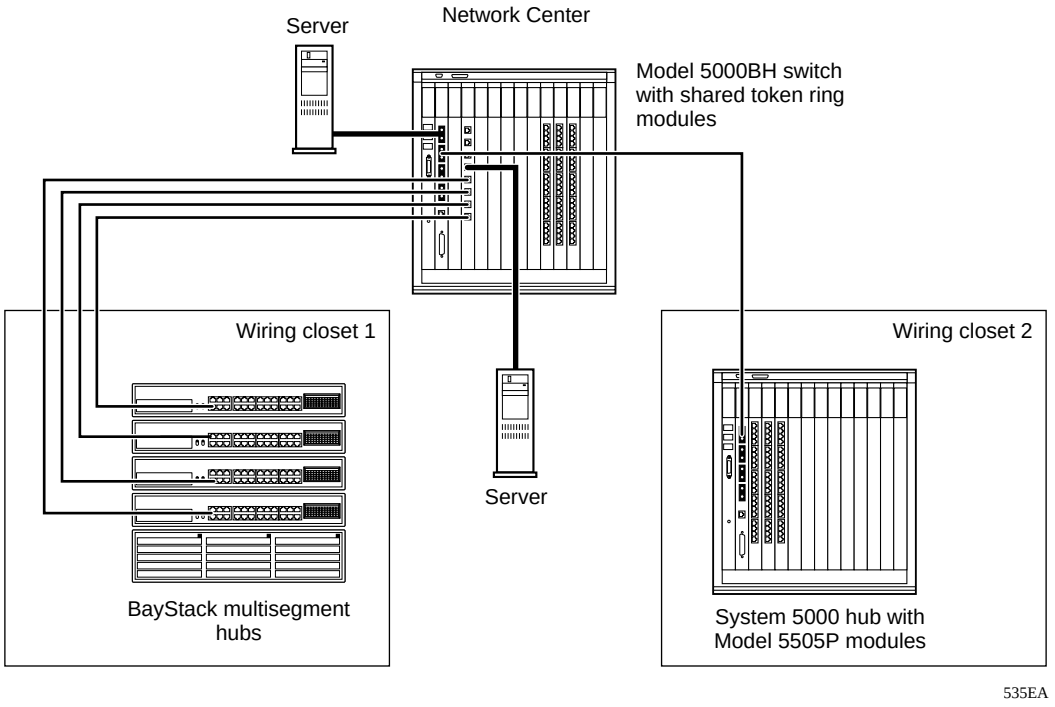

Figure 2-8. Collapsing All Ring Segments to a High-Density Switch

The example shown in Figure 2-8 collapses the multiple ring segments in the BayStack™ and System 5000™ hubs into the high-density Model 5000BH switch. Instead of having a single ring as the backbone, the Centillion backplane allows for higher bandwidth at lower cost.

Using Distributed Backbone Switching Without LANE

Figure 2-9 shows a network with a Model 5000BH switch in the network center and two separate wiring closets. The Model 5000BH switch provides distributed backbone switching between the multiple rings within the two closets.

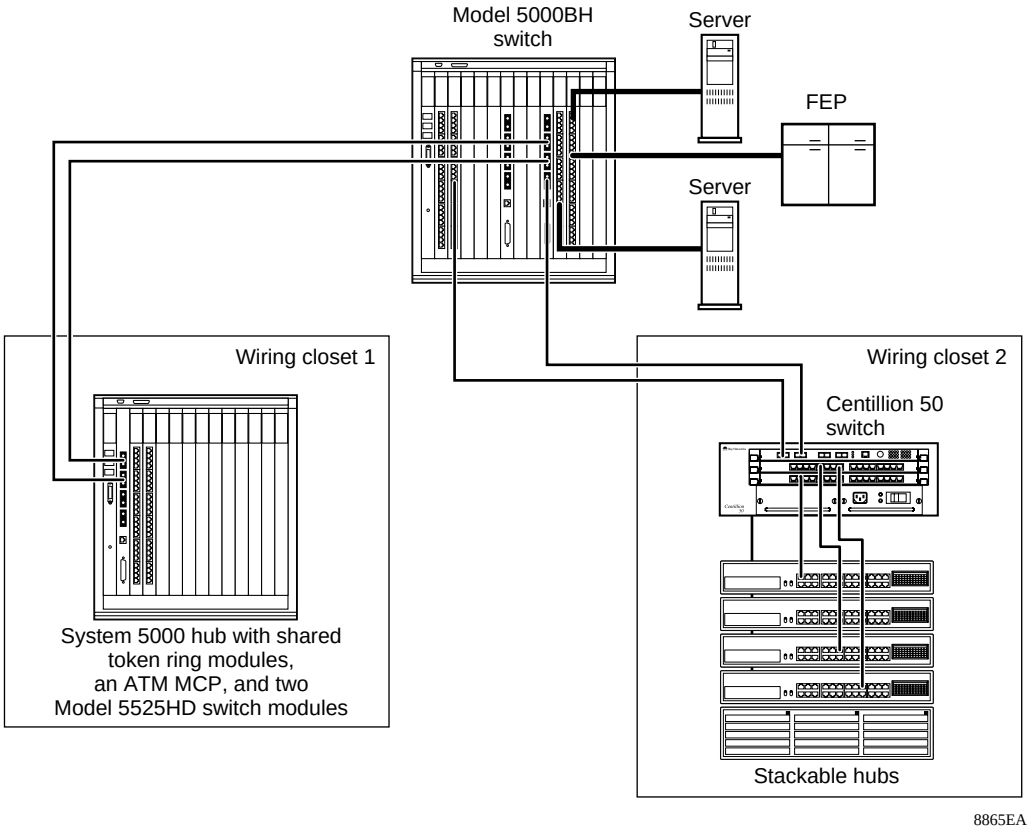

Figure 2-9. Distributed Backbone Switching Without LANE

Wiring closet 1 uses a System 5000 hub to create an integrated shared/switched network. The Model 5525HD switch module allows switching between backplane ring segments and front RJ-45 attached devices on ports 4 through 8. These ports have three modes to choose from:

- Use any of the RJ-45 ports on the front of the module as a switch port to direct-attached devices.
- Use ports 4 through 8 to switch between backplane rings 1 through 5, respectively.
- Combine the two methods to allow switching between shared backplane rings as well as providing dedicated switch ports out of the front of the module. Devices sharing a switch port with the backplane ring are part of the same ring.

If your existing backbone has a System 5000 hub, then you can install a Model 5525HD switch module to provide a distributed backbone for the shared backplane rings.

Wiring closet 2 uses a Centillion 50/100 switch with TokenSpeed modules to switch between stackable token ring hubs. The System 5000 hub and the Centillion 50/100 switch both use an ATM MCP module to connect to the Model 5000BH switch for switching between the multiple rings within each wiring closet.

The Model 5000BH switch in the network center also uses a token ring switch module to provide direct attachments for servers and a front end processor (FEP). A Centillion 50/100 switch could be used in the network center to provide the same switching solutions.

Using Distributed Backbone Switching with LANE

Figure 2-10 shows a network that has moved its central services to ATM connections. The Model 5000BH switch can provide LANE services over IISP and/or PNNI links to LAN emulation clients on the System 5000 hub and the Centillion 50/100 switch.

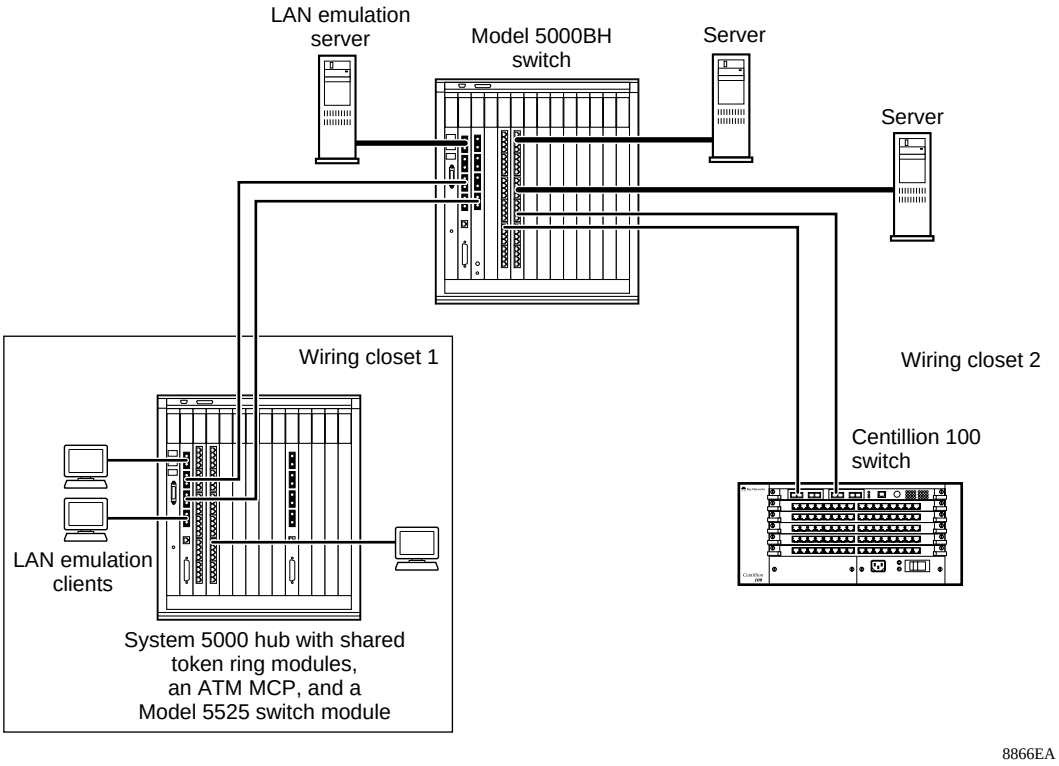

Figure 2-10. Distributed Backbone Switching with LANE

Ethernet Solutions

This section provides solutions for the following Ethernet network needs:

- Giving a server a dedicated port (this page)
- Segmenting an Ethernet LAN (this page)
- Replacing bridges or routers (page 2-15)
- Concentrating LANs into a router (page 2-16)
- Replacing a LAN backbone (page 2-18)

Giving a Server a Dedicated Port

Giving a server a dedicated port on the switch improves server response and increases throughput capacity by containing broadcasts and lessening collisions.

Segmenting an Ethernet LAN

Excessive traffic on an Ethernet LAN slows response time to all stations. Using the switch, you can segment a congested LAN to create multiple LANs, thereby increasing the average bandwidth available to each device.

The example in Figure 2-11 gives users 1 through 6 dedicated Ethernet ports for improved performance. The rest of the users are connected through a hub. This example also applies to the Centillion 50 or Model 5000BH switch.

Before

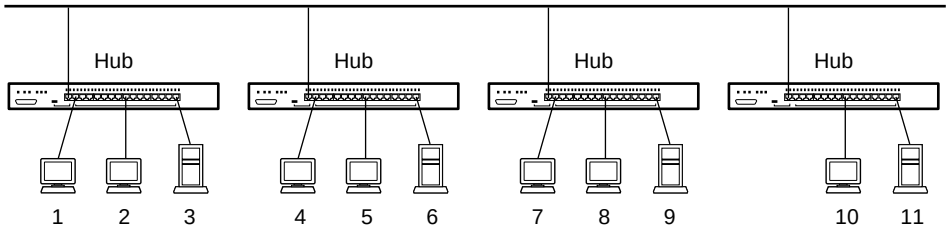

After

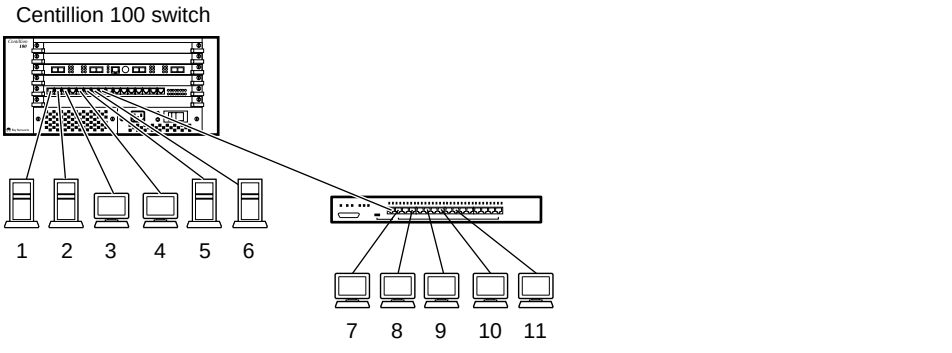

646EB

Figure 2-11. Segmenting an Ethernet LAN

Special consideration is required for enabling spanning tree on a port with a dedicated end station attached to it. Enabling spanning tree on a port dedicated to an end station client will delay network connectivity to that server if its interface is restarted or reconnected to that port. The spanning tree algorithm takes about 30 seconds to go from the Blocked, Listening, Learning states to the Forwarding state. When spanning tree is disabled for a port with a dedicated end station attached, that port will not participate in the spanning tree algorithm and will be placed immediately in the forwarding state.

Replacing Bridges or Routers

The switch offers two benefits as a replacement for bridges and routers. First, it provides additional ports at lower equipment cost. Second, it increases the aggregate bandwidth of the LAN backbone.

Figure 2-12 shows a Centillion 100 switch replacing an Ethernet bridge. This example also applies to the Centillion 50 or Model 5000BH switch.

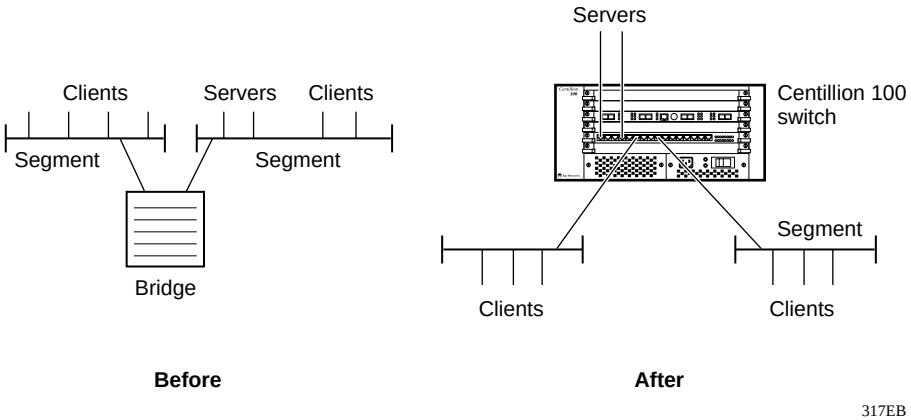

Figure 2-12. Replacing an Ethernet Bridge

The example in Figure 2-12 gives the servers dedicated ports. The clients are connected to a hub, and the hub is connected to the switch.

Concentrating LANs into a Router

The switch can relieve an overloaded router of port and processing burdens. The router in the network illustrated in Figure 2-13 has no capacity for additional Ethernet segments, which prohibits the network manager from segmenting further.

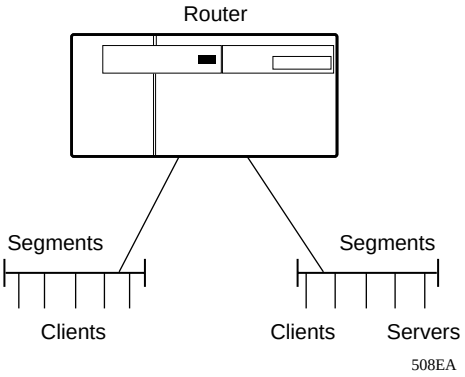

Figure 2-13. Ethernet Network with No Growth Capacity

Figure 2-14 uses the switch to add segments and to replace many Ethernet ports on the router with a single ATM port, which can route among many subnets. This configuration allows the router to route among multiple VLANs on a single physical link that has multiple logical links configured on it.

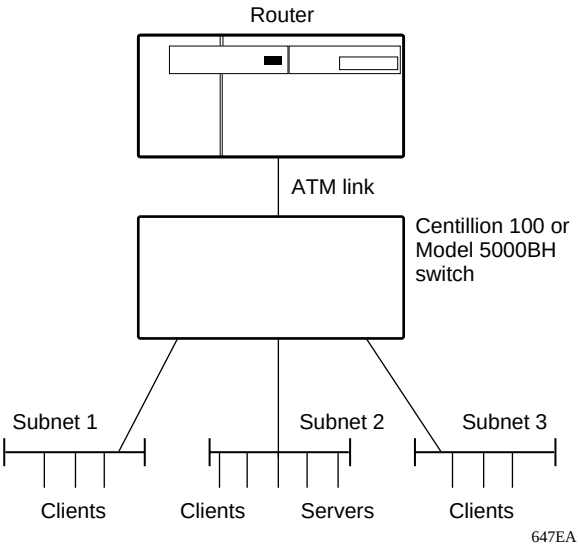

Figure 2-14. Switch Concentrating Multiple Segments as a Single Bridged Subnet

Replacing a LAN Backbone

Backbones can become bottlenecks as traffic between LAN segments increases. The switch can serve as a collapsed backbone, increasing the aggregate bandwidth of the backbone to multiple gigabits per second.

Figure 2-15 shows an existing Ethernet backbone using routers and bridges.

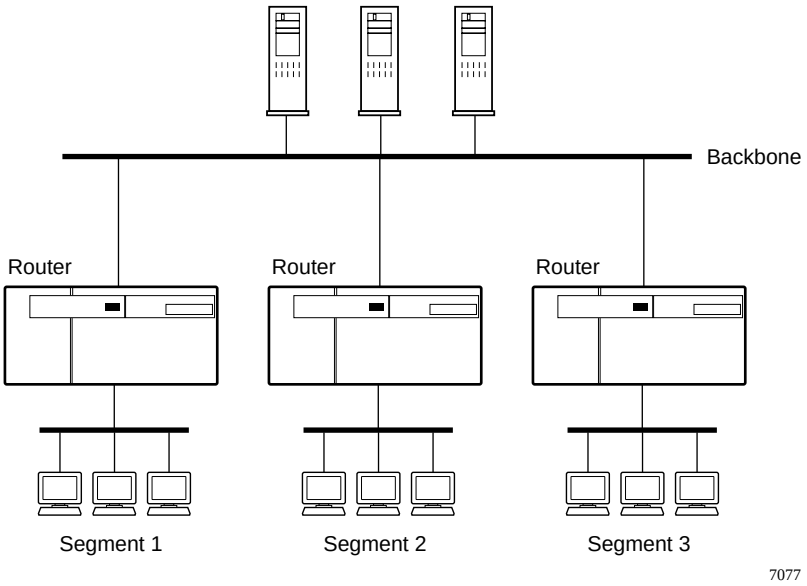

Figure 2-15. Ethernet Backbone Using Routers and Bridges

Figure 2-16 shows an example of increasing Ethernet backbone throughput with a Centillion 100 switch. A Centillion 50 or Model 5000BH switch could be substituted to provide the same throughput improvements.

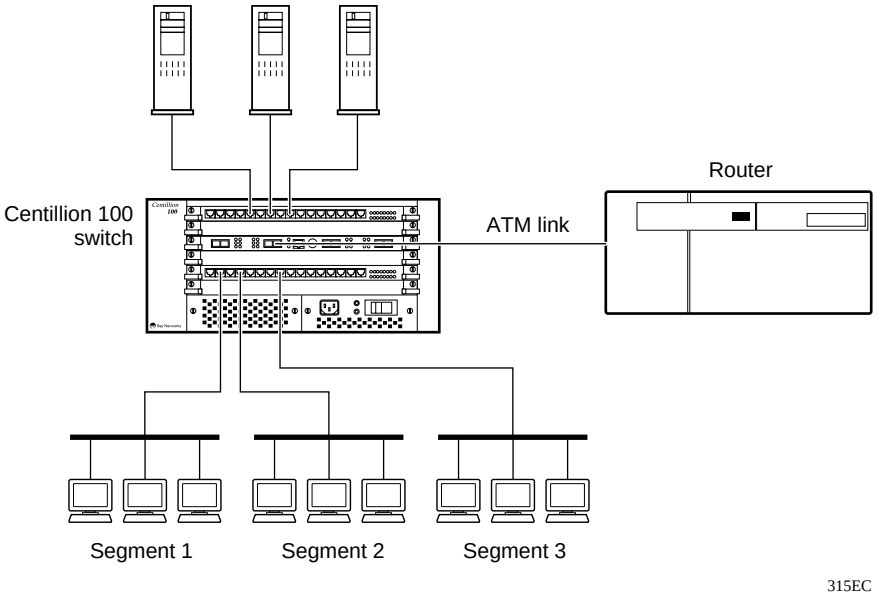

Figure 2-16. Increasing Ethernet Backbone Throughput with a Centillion 100 Switch

Chapter 3
MPOA in an ATM Network

Multi-Protocol Over ATM (MPOA) is the ATM Forum standard that provides distributed layer 3 forwarding capabilities for LAN traffic across an ATM network without requiring the traffic to pass through a multiprotocol router. The physical separation of internetwork layer route calculation and forwarding is called virtual routing. For a detailed explanation of MPOA, refer to the ATM Forum specification *Multi-Protocol Over ATM, Version 1.0* (AF-MPOA-0087_000).

This chapter includes the following sections:

- MPOA basics (page 3-2)
- MPOA components (page 3-3)
- Establishing a shortcut (page 3-7)

MPOA allows large ATM networks to establish virtual circuits between edge devices. In conjunction with the Centillion 50/100 and Model 5000BH switches, MPOA is implemented in the Backbone Node (BN®) router with an ATM routing engine (ARE) and the Centillion Multiprotocol Engine (MPE).

Note: A switch must be configured for LUNI 2.0 in order to use MPOA. You configure LUNI support in the ATM Configuration windows in SpeedView.

Virtual routing provides the following benefits:

- Decreased number of devices that must be configured to perform internetwork layer route calculations
- Increased scalability by increasing the number of devices participating in internetwork layer route calculations
- Reduced latency by reducing the number of points where packet processing occurs (Traffic is forwarded over an ATM VC incurring only a single network hop and incurring a single segmentation and single reassembly.)

MPOA Basics

As network traffic increases, a router can become a bottleneck for layer 3 data forwarding. MPOA increases the layer 3 capabilities of your network by allowing MPOA clients in the LANE network to forward layer 3 traffic. The router is still used for its forwarding and topology information, but MPOA offloads a large amount of the forwarding responsibilities to the MPOA client.

Having increased forwarding capabilities means that you can use your existing layer 2 equipment to perform layer 3 functions and improve overall performance in your network.

MPOA is comprised of two main components:

- MPOA clients (MPCs)
- MPOA servers (MPSs)

The clients and servers interact to determine the need for virtual circuit connection (VCC) shortcuts between edge devices and to resolve the addresses of the edge devices. This address resolution occurs through a combination of layer 2 and layer 3 translation and communication using the Next Hop Server (NHS) and Next Hop Routing Protocol (NHRP).

MPOA Components

An MPOA network includes MPOA clients and MPOA servers. Next Hop Route Protocol is used by MPSs that are logically separated by more than a single router to resolve ATM addresses. MPCs reside in the Centillion 50/100 and Model 5000BH switches; MPSs may reside in the BN ATM routing engine and Centillion MPE.

The MPOA client process is associated with one or more LAN emulation clients. This process occurs in an edge device or LAN emulation host and monitors layer 3 traffic. It requests and initiates MPOA VCC shortcuts and receives inbound VCC shortcuts. The MPC performs internetwork layer forwarding, but it does not run internetwork layer routing protocols.

The MPOA server coexists with a router and interfaces with LANE for layer 2 information. It communicates with other MPOA devices to resolve layer 3 addresses to ATM addresses using NHRP and MPC cache imposition requests.

MPOA Clients

MPCs determine the need for and initiate internetwork shortcuts, as well as perform internetwork layer forwarding.

When an ingress MPC determines that a data flow meets the requirements for a shortcut, it sends out an MPOA resolution request for the information needed to create a shortcut to the destination address.

When an egress MPC receives data from an ingress MPC across a shortcut, the MPC receiving the inbound shortcut traffic forwards the traffic to the higher layers (as in the case of an ATM attached host) or adds the appropriate data link layer (DLL) encapsulation and forwards the traffic out of the MPC (as in the case of an ATM edge switch port on a Centillion 100 or Model 5000BH switch) to the destination station.

On the Centillion 50/100 and Model 5000BH switches, a single MPC supports all LECs running on that edge device. MPOA must be enabled by the network administrator on a per LEC basis. The MPC can receive its configuration information from an LE_Configure Request sent to the LAN emulation configuration server by the LEC. The MPC can also store configuration parameters locally. An MPC device type length value field (TLV) in the LE_ARP signals that MPOA is enabled.

Note: By default, MPCs receive their configurations from the LECS. You can create a remote (on the LECS) or local MPOA client configuration using SpeedView.

Conceptually, an MPC has two subsystems: control and data. Control subsystems send and process control messages, such as MPOA resolution requests. Data subsystems send and receive data traffic. The ATM address associated with the control subsystem is referred to as a control ATM address. The control ATM address is associated with the MCP switch module. An ATM address associated with a data subsystem is referred to as a data ATM address. Data ATM addresses are associated with host switch modules. Both control and data ATM addresses are generated automatically at system startup.

Ingress MPC Functionality

An ingress MPC detects flows of packets forwarded over an ELAN to a router that contains an MPS. The ingress MPC learns the MAC address of MPSs located in the same ELAN through a device type TLV in the LE_ARP response. When the MPC recognizes that a flow could benefit from a shortcut (when it recognizes data traffic bound for the same destination IP and router MAC address at a user-configured rate) the MPC looks in the ingress cache. If the MPC finds the destination address in its local cache, it immediately begins to establish the SVC. If the MPC fails to locate the destination address in the ingress cache, it generates an MPOA resolution request to the appropriate MPS.

Egress MPC Functionality

An egress MPC receives internetwork data frames from other MPCs to be forwarded to its local users. When an egress MPC receives a frame over a shortcut, it adds the appropriate DLL encapsulation and sends it on the bridge ports or internal host stack. The MPC receives the DLL encapsulation information from the egress MPS. DLL encapsulation information is stored in the egress cache of the MPC.

The egress MPC sends a cache imposition reply to the egress MPS for each cache imposition request it receives. The egress MPC must first determine if it has enough resources to maintain the cache entry and receive a new VCC. If resources are not available, error status information is inserted into the reply.

MPOA Servers

The MPOA server provides internetwork layer forwarding information to MPCs. One or more MPSs exist on a router, such as an ATM router engine or a Centillion Multiprotocol Engine. The MPS uses NHRP to get routing information, to answer MPOA resolution requests from ingress MPCs, and to provide DLL encapsulation information to egress MPCs.

Figure 3-1 illustrates how the MPS interacts with the Next Hop Server (NHS), the associated LECs, and MPOA clients.

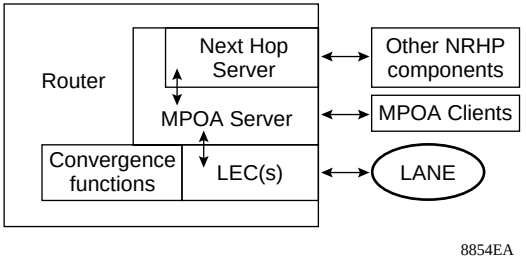

Figure 3-1. MPOA Server and Components

Ingress MPS Functionality

The ingress MPS processes incoming MPOA resolution requests from MPCs. If the request is for a local destination, the MPS answers with an MPOA resolution reply locally without contacting another MPS. If the request is not for a local destination, the MPS must translate the MPOA resolution request to an NHRP resolution request and send it to the next MPS.

Note: If there is only one MPS in the network, then there is no MPOA to NHRP translation.

Egress MPS Functionality

When an NHRP resolution request is received, the egress MPS translates the NHRP resolution request to an MPOA cache imposition request and sends it to the egress MPC. The cache imposition request provides encapsulation and state maintenance information to the egress MPC.

After sending the cache imposition request, egress MPS waits for state, address, and ingress tagging information in an MPOA cache imposition reply from the egress MPC. Then, the egress MPS formulates an NHRP resolution reply and sends the information back to the original requestor (the ingress MPS).

If resources are available, the reply provides the ATM address and success status.

Discovery

One MPC represents multiple LECs. MPOA devices configure and communicate their device type (MPC, MPS, or both) and their ATM address to the clients they serve. LECs advertise their associated MPOA devices using a type, length, and value (TLV) variable in the LE_REGISTER and in all LE_ARP requests and replies.

Establishing a Shortcut

Figure 3-2 shows how data packets flow with and without MPOA. When a data packet reaches an ingress MPC (MPC 1) in an MPOA system, by default, the packet is forwarded over LANE to the router.

However, if the MPC recognizes the same destination IP address going to the same router MAC address and the rate of flow reaches a user-configured rate (such as 10 packets per second), the MPC determines that the flow could benefit from a shortcut in order to bypass the router. The MPC sends an MPOA resolution request to the MPS for the ATM address of a virtual circuit to the destination IP address (MPC 2).

The MPS responds with the ATM address of the egress MPC, and a unicast, unidirectional shortcut is established. After the shortcut is established, the packet is stripped of its DLL encapsulation, encapsulated in the tagged MPOA format, and forwarded to the egress MPC over the shortcut VCC.

Note: MPOA shortcuts are for unicast, unidirectional communication only. An existing VC between MPC 1 and MPC 2 may be reused for a different connection, but only after the MPOA resolution process is completed for that new connection. MPOA shortcuts do not carry broadcast or multicast traffic.

The egress MPC receives the packet and attaches the DLL header. The frame is passed to the egress station as though it had arrived directly from the LEC of the egress MPS.

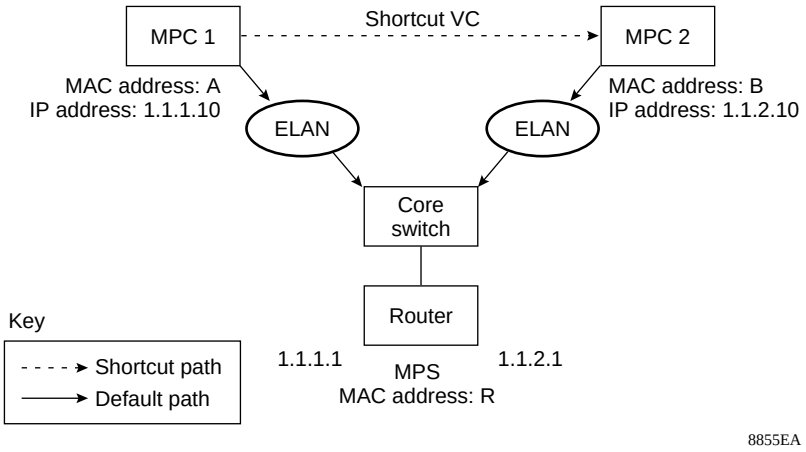

Figure 3-2. Example of Data Pack Flow with and Without MPOA

Figure 3-3 describes the MPOA resolution process for establishing a shortcut.

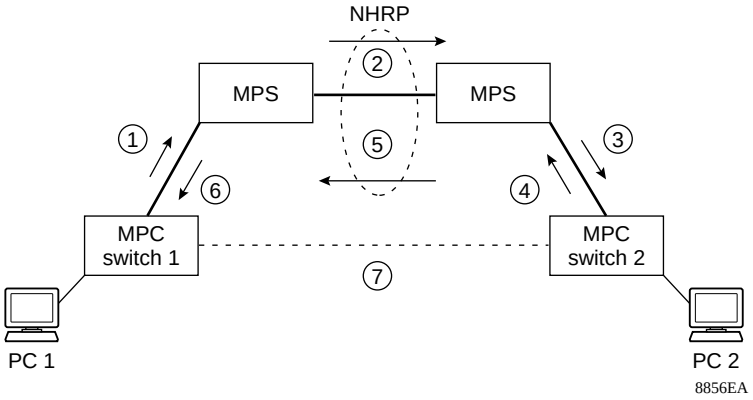

Figure 3-3. Establishing a Shortcut

In Figure 3-3 the MPCs and MPSs go through the following steps to establish a shortcut. PC 1 is sending a packet to PC 2. The MPC on switch 1 recognizes that this connection can benefit from a shortcut.

1. The ingress (source) MPC in switch 1 sends an MPOA resolution request to the ingress MPS.
2. The MPS translates the MPOA resolution request to an NHRP resolution request and sends it to the target (egress) MPS. (If there is only one MPS in the network, then this step does not occur.)
3. The egress MPS translates the NHRP resolution request to an MPOA cache imposition request and adds a MAC header (a DLL) and state information.
4. The egress MPC receives the MPOA cache imposition request and sends an MPOA cache imposition reply and the ATM address of the EtherSpeed™ module to which PC 2 is attached back to the egress MPS.
5. The egress MPS translates the MPOA cache imposition reply into an NHRP resolution reply. (If there is only one MPS in the network, then this step does not occur.)
6. The ingress MPS translates the NHRP resolution reply into an MPOA resolution reply that includes the ATM address of the EtherSpeed module to which PC 2 is attached.
7. Using the ATM address provided in the MPOA resolution reply, switch 1 calls switch 2 to set up a VCC that carries the MPOA shortcut.

You can also create internal shortcuts for clients within the same switch (or same module) but on different ELANs. Figure 3-4 shows how you can create a connection between PC 1 and PC 2 without having to leave the switch.

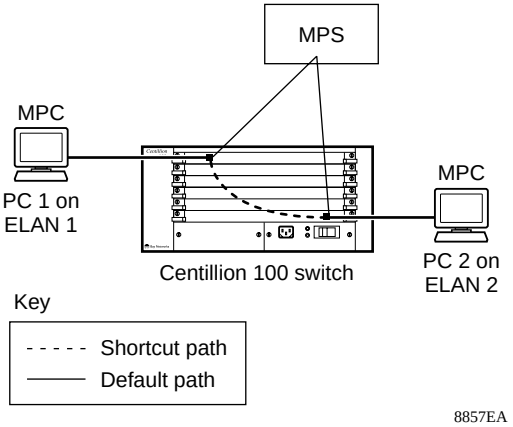

Figure 3-4. Internal MPOA Shortcut

Virtual Circuit Multiplexing

When MPC 2 sends an IP response after receiving the data packet in Figure 3-2 on page 3-8, it does not use the same shortcut automatically. MPC 2 must go through the same recognition and address resolution process that MPC 1 did in order to establish a shortcut. After the same route is discovered by MPC 2, it will send a resolution request to the MPS; if a shortcut already exists between the two points, MPC 2 can share that VC.

Figure 3-5 provides an example of multiplexing virtual circuits between two Centillion switches.

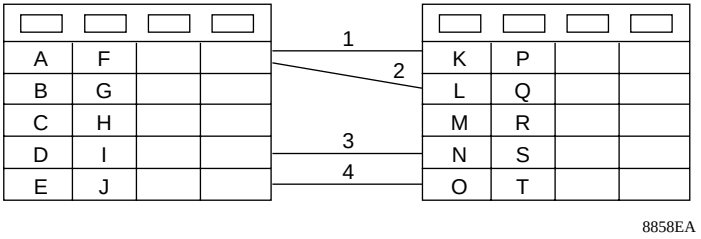

Figure 3-5. Multiplexing Virtual Circuits

The letters A through T represent 20 LECs that are representative of their respective Ethernet ports (A through T). Client A represents an IP client in the IP subnet A that is attached to the EtherSpeed port A.

If an MPOA flow is detected between clients A and K, the MPOA resolution process occurs, and a shortcut will be set up as SVC 1. All unicast IP traffic that is destined from A to K is sent across SVC 1. However, all unicast IP traffic that is destined from K to A will continue to be sent across the default path through the default gateway of client K until the MPC that represents K detects a flow and initiates the address resolution process.

After the flow has been detected from K to A, client K begins sharing SVC 1. No new SVC is set up. To take this one step further, suppose that a flow is detected from P to F. Because each slot in the switch is seen as one unique ATM address, client P recognizes that there is a circuit already established between the two switches. Therefore, P can also use SVC 1.

Table 3-1 represents the SVCs that can be used for the various flows possible in Figure 3-5.

Table 3-1. Possible SVC flows

From	To	SVC
A or F	K or P	1
A or F	L or Q	2
D or I	N or S	3
E or J	O or T	4

Chapter 4
PNNI in an ATM Network

This chapter describes the Bay Networks implementation of the ATM Forum Private-Network-to-Network Interface (PNNI) protocol. PNNI is based on the ATM Forum PNNI 1.0 specification. Information is also included about using PNNI along with IISP in a network. For more information about using IISP, refer to Chapter 5, "IISP in an ATM Network."

This chapter discusses the following topics:

- What is PNNI (this page)
- Using PNNI in the ATM network (page 4-2)
- Using PNNI with Centillion 50/100 or Model 5000BH switches (page 4-3)
- Using PNNI with IISP (page 4-9)
- Configuring PNNI (page 4-14)
- PNNI features (page 4-18)

What Is PNNI

PNNI includes a routing protocol for distributing network topology information among switches and clusters of switches and a signaling protocol for defining message flows to establish point-to-point or point-to-multipoint connections. PNNI routing protocol is a dynamic routing protocol that enables the ATM switches to learn about other ATM switches in the network dynamically so that any link or switch failures can be detected and the calls rerouted. The network can have any number of parallel/redundant paths without users worrying about loops in the network because an ATM switch can use the dynamically learned network topology to generate a loop-free source route.

Because PNNI is an ATM Forum standard, it allows for interoperability of switches from different vendors and designs that span multiple organizations.

In IISP, every ATM device that might need to make a choice about the routing of ATM calls must have the choices statically listed in a routing table; the only way to change the routing table is by manually attaching the device to the switch and modifying the routing table. PNNI allows switches to learn about the network dynamically and exchange routing databases, enabling them to automatically respond to changes in resource availability. Database summary information includes the exchange of link states between switches. The link group feature used in IISP networks to prevent loops in the network is not required in PNNI networks.

The PNNI signaling protocol enables multiple signaling channels on a single physical port through the use of multiple virtual path identifier (VPI) values (see "Logical Links" on page 4-19).

Using PNNI in the ATM Network

PNNI resolves data paths for source routed call requests using source routing across an ATM network. PNNI's dynamic nature means that link failures or changing path characteristics impact the route of call requests. Rerouting can occur while simultaneously accounting for bandwidth, QoS, and other route requests. Parallel PNNI links between nodes support call-balanced, alternate-path redundancy.

PNNI switches form themselves into peer groups based on their network service access point (NSAP) address. If the NSAP addresses are the same up to a certain point (called the node level), the switches will be part of the same peer group. From the outside of a peer group, the entire group is represented as a single node (the logical node). Within a peer group, all nodes are peers and must see the full topology of that peer group.

PNNI group members find each other by exchanging hello packets on their NNI links. When they have established that the other end of the link is in the same peer group, they send PNNI topology state packets (PTSPs) to exchange routing information. The PTSPs contain one or more PNNI topology state elements (PTSEs) for a particular node. Each PNNI node advertises its resources and the outbound resources of each attached link. Logical nodes communicate with each

other using logical links. On a logical (not physically attached) link, hello sending is done by way of a routing control channel (RCC) based on a switched virtual channel (SVC) connection. PNNI has been designed specifically for use in routing SVCs across an ATM network.

Within each peer group, one switch, such as a Centillion 1000 switch, will be elected peer group leader. Neighbors that are in different peer groups will not exchange topology information and are “border nodes.” However, PNNI switches may also be configured as a single PNNI group with no hierarchy, where it is not necessary to elect a group leader.

One use of PNNI is in the ATM LAN backbone. Typical configurations connect frame-based desktop systems to an ATM edge switch, which connects to the ATM backbone through a core switch. ATM servers and routers can connect directly to the ATM backbone through the core switch or the edge switch. A big advantage of ATM backbones over frame-based backbones is topology knowledge. ATM nodes supporting PNNI understand the ATM topology and provide traffic flow protection against link failures by establishing load balancing across redundant ATM links. In contrast, UNI connections with edge switches do not participate in network topology.

As with IISP, PNNI provides network loop protection and redundant links as described in Chapter 5, “IISP in an ATM Network,” in the sections on “How a Link Group Provides Loop Prevention” on page 5-7 and “Example of IISP and Redundant Links in a Network” on page 5-11. The difference is that these functions are done automatically with PNNI.

Using PNNI with the Centillion 50/100 or Model 5000BH Switches

The advanced image for the Centillion 50/100 and Model 5000BH switches supports PNNI. You can use SpeedView to configure PNNI within each switch.

The hierarchical network design shown in Figure 4-1 is an example of a PNNI network using Centillion 50/100 or Model 5000BH switches. This network includes two core switches and a number of transit switches and allows for a PNNI hierarchy of switches with each transit switch pair and its associated edges forming the lowest-level PNNI group. Table 4-1 shows a possible addressing scheme for this network.

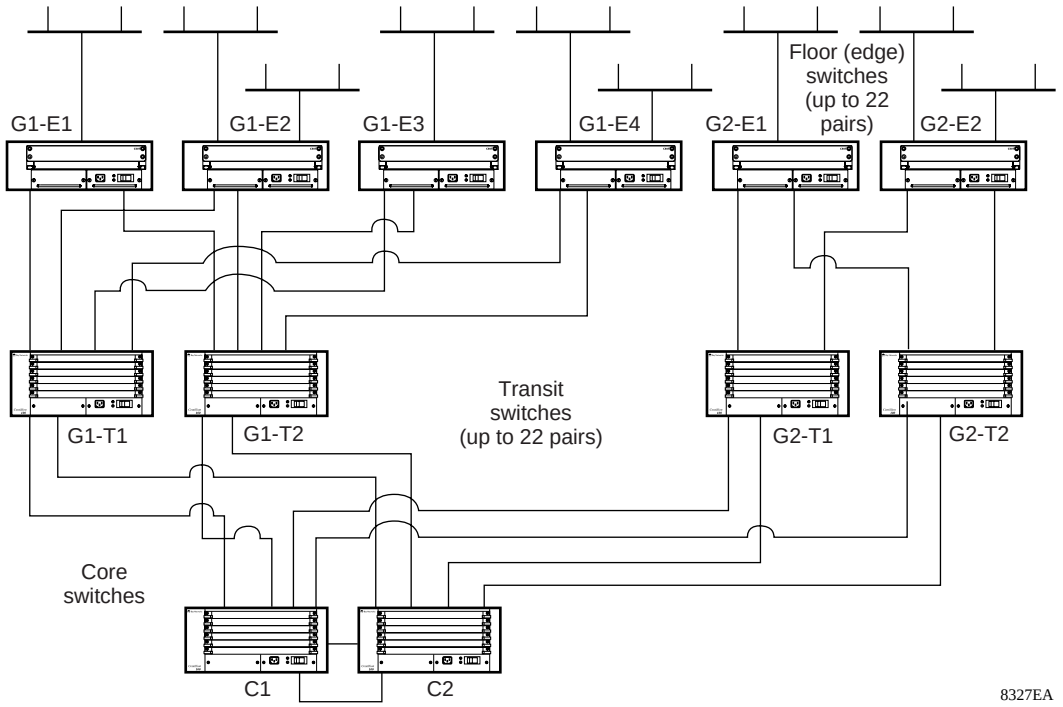

Figure 4-1. PNNI Network Using Centillion 50/100 and Model 5000BH Switches

Table 4-1. Addressing Scheme for a PNNI Network

Peer Group/Switch	PNNI Address	Level
Core	39.11.22.33.44.55.66.77.00	72
C1	39.11.22.33.44.55.66.77.00.00.00.00.01	
C2	39.11.22.33.44.55.66.77.00.00.00.00.02	
Group 1	39.11.22.33.44.55.66.77.01	72
G1-T1 (transit switch)	39.11.22.33.44.55.66.77.01.00.00.00.01	
G1-T2 (transit switch)	39.11.22.33.44.55.66.77.01.00.00.00.02	
G1-E1	39.11.22.33.44.55.66.77.01.00.00.00.03	
G1-E2	39.11.22.33.44.55.66.77.01.00.00.00.04	
G1-E3	39.11.22.33.44.55.66.77.01.00.00.00.05	
G1-E4	39.11.22.33.44.55.66.77.01.00.00.00.06	

Table 4-1. Addressing Scheme for a PNNI Network (continued)

Peer Group/Switch	PNNI Address	Level
Group 2	39.11.22.33.44.55.66.77.02	72
G2-T1 (transit switch)	39.11.22.33.44.55.66.77.02.00.00.00.01	
G2-T2 (transit switch)	39.11.22.33.44.55.66.77.02.00.00.00.02	
G2-E1	39.11.22.33.44.55.66.77.02.00.00.00.03	
G2-E2	39.11.22.33.44.55.66.77.02.00.00.00.04	
Group 3	39.11.22.33.44.55.66.77.03	72
.		
.		
.		

A single PNNI group with Centillion 50/100 and Model 5000BH switches should be no larger than approximately 32 switches to limit the size of the routing table and maximize the speed of PNNI reconvergence. However, because the switch can perform address summarization and propagation of external (IISP) routes, you can link PNNI groups of switches together. The size of the PNNI group does not limit the maximum size of the network.

Although PNNI is a system to make management of ATM devices simpler, as with all other networking issues, a good design is important to keep a network resilient and responsive. The most important rule for PNNI is to keep the routing tables simple; doing so allows quicker recovery when there is a failure. It is important to ensure that the combination of IISP and PNNI does not lead to a situation where the failure of a device prevents some calls from being established. It is also important to use optimum paths to route calls between groups. In the future, with universal PNNI, it will be important that addresses and peer groups be efficiently organized so that PNNI routing tables remain manageable, no matter how many switches are in a single PNNI hierarchy.

The simplest rule to follow is to ensure that a set of switches in a particular location (a set of switches that have the same PL byte) has a simple hierarchy with all PNNI switch groups connecting directly to the core group. Therefore, in a campus environment, rather than making each building its own PNNI hierarchy, the core switch group should be represented in each building because the transit switches in each building do not need to make a decision about where to route traffic—they can just route it to the core, which is the same whether or not a link fails—and the core switches have enough knowledge of the rest of the core that they can intelligently route around a failure. This issue is particularly relevant in a

network that does not have Centillion 50/100 or Model 5000BH switches to form a hierarchy. In this case, it is likely that IISP would be used to interconnect PNNI peer groups and failure information might not be properly disseminated. Figure 4-2 shows an example of a good network design using PNNI and a design using IISP links that might cause a difficult failure situation.

In the good design, the IISP links are inside the buildings and there is a PNNI group spanning the campus that includes all the interbuilding links. If one of these interbuilding links fails, the PNNI protocol will inform the other switches in the peer group about the failure and calls will be routed around the failure. In the ineffective design, if a link fails between two buildings, the other building peer groups will not be aware of the failure and may not make a correct routing decision when attempting to set up an SVC.

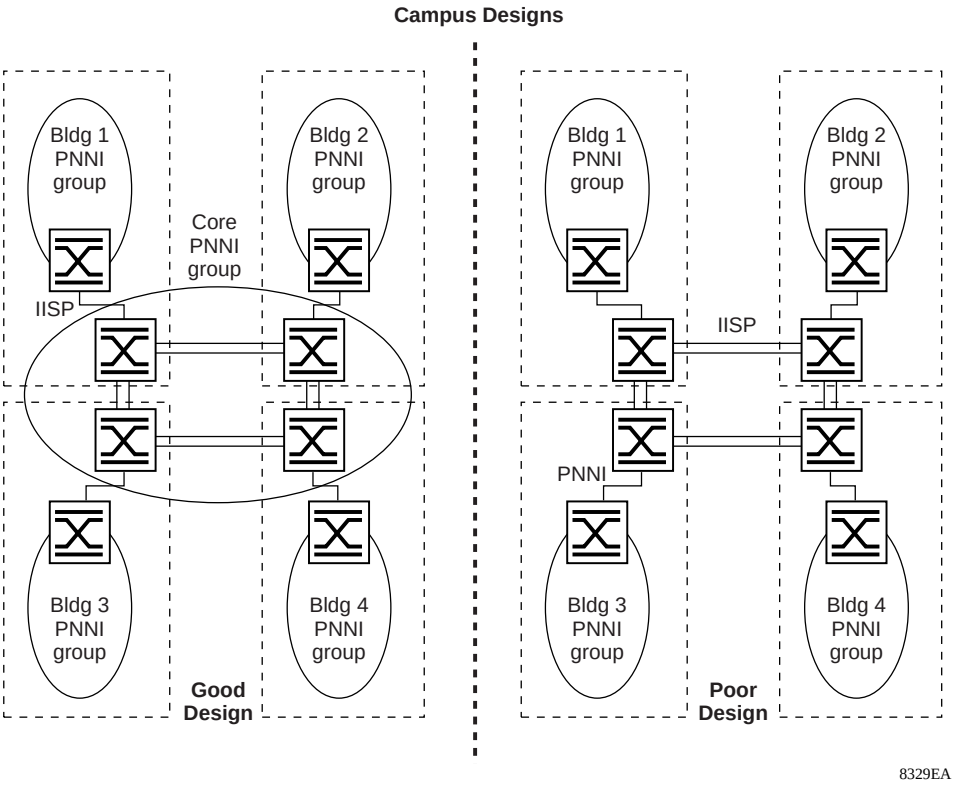

Figure 4-2. Good and Inadequate PNNI Network Designs

On the other hand, hierarchical PNNI does allow any more flexibility in terms of the sorts of campus and core connections that are possible. Although each of the four examples in Figure 4-3 are valid uses of PNNI, if the designs are used in a pure IISP or mixed IISP/PNNI environment problems might occur.

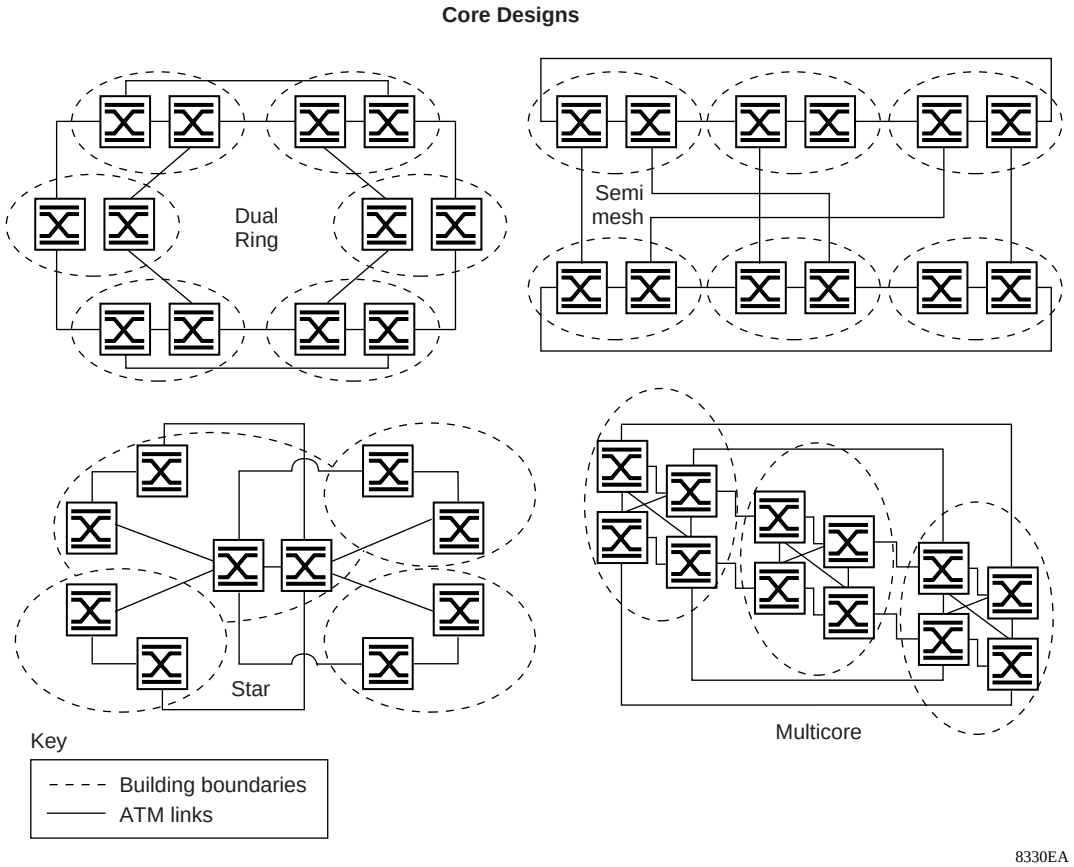

Figure 4-3. Sample Core Designs

All the designs show possible ways to connect a number of buildings in a campus. The design you choose depends on the physical availability of fiber and on the expected data flows between the buildings.

The dual ring and multicore designs are best suited for campus interconnects where the primary data traffic is local to a building but there is some need to connect across the campus. The two designs can be combined so that some buildings have the meshed core of four switches and others have only two switches.

The star design is the obvious preferred design for a single core data center and a number of outlying buildings. This design is preferred when data flow is primarily hierarchical, but physical constraints such as the distance between buildings and the numbers of available fiber runs may not allow this design to work.

The semimesh design is the type of network design to use when there is significant intracampus data flow among various buildings. As depicted, both this design and the dual ring design allow the six buildings to connect and use the same number of links and fiber runs out of each building. By redistributing the cabling, the semimesh design provides better connectivity than the dual ring design. However, physical constraints (for example, distance and numbers of fiber runs) may not allow this design to work.

You can use these designs in three ways:

- Each building can be a PNNI peer group with IISP links between them. In a pure Centillion 50/100 and Model 5000BH environment, this arrangement will not provide the necessary fault tolerance for some of the designs.
- You can have a campus PNNI peer group with switches in each building as shown in the diagrams. Building switches are connected to these via IISP.
- The switches can be Centillion 50/100 and Model 5000BH, and the buildings can be set up as a hierarchy. You can attach additional ATM switches to the Centillion 50/100 and Model 5000BH switches via PNNI or IISP. Note that any Centillion 50/100 or Model 5000BH switches must use IISP to get to this hierarchical core.

The dual ring is another version of an inefficient design. In a hierarchical PNNI environment or in one where, as in the example, a network with 12 switches is a single peer group, it works well. However, if the individual buildings are each peer groups that are connected via IISP, then the network will not be fault tolerant if a parallel pair of links fails.

The multicore design can be treated like the dual ring in many ways. However, notice that the design drawn with only three cores does not have the recovery problems of a ring with more than three cores.

The semimesh design is the most suitable if you understand that, as with the dual ring design, it is possible that the recovery time for some link failures might not be as fast as they could be. Because this design works well for sites that have a heterogeneous traffic mix, this type of failure might not impact the majority of users.

The star design is suitable for all configurations—the important factor for this design is to be sure to that fiber runs from any building to the core building are placed in different conduits to avoid loss of connectivity if one of the fiber runs is severed.

Using PNNI with IISP

Although it would be desirable to have PNNI throughout an entire ATM network, there may be cases where PNNI must run in conjunction with IISP in a network to accommodate third-party switches or previous software versions. IISP routes can be propagated through a PNNI domain, and IISP can be used to interconnect nonhierarchical PNNI groups. However, there are two caveats when mixing PNNI and IISP:

- If a link is configured as PNNI, it cannot also be IISP; if it is configured as IISP, it cannot also be PNNI. Therefore, switches at both ends of the link must have the link configured the same way (IISP or PNNI).
- If it is not possible to set up a call as originally intended, PNNI can perform rerouting, but this capability does not extend across IISP automatically. It is important to always have at least two IISP links outside of the switch to allow for rerouting. Refer to page 4-13 for examples where this requirement has not been applied.

PNNI routing tables can get very large if the network topology is complicated and there are multiple routes between switches. It is possible for PNNI to do route summarization to reduce the number of entries in the tables, but this approach works only if the addressing scheme has been carefully thought out and implemented. In a complicated topology, route reconvergence after a failure could be very time consuming.

One good reason for mixing PNNI and IISP is to prevent routing tables from becoming unwieldy. It also prevents unnecessary propagation of topology changes that are not relevant to distant switches. Another reason for mixing PNNI and IISP is that PNNI enables the switch to do more intelligent call balancing across multiple links, which should result in greater efficiency and improved behavior during recovery from a failure.

One configuration is to use IISP for the links between the transit switches and the core. If the ATM network is expected to expand, either in terms of the number of switches in a location or by interconnecting remote locations, it is best if the core is a PNNI group even if it consists of only a pair of switches. This configuration ensures that, as the network expands, the correct addressing will be used.

For small and medium-size networks, it will be unnecessary to have PNNI at the edge and PNNI in the core; and, in this case, either the edge (wiring closet) or the core (data center) should be solely IISP. There are advantages and disadvantages to both of the approaches.

Using PNNI in the edges or wiring closets with IISP in the core or data center has the advantage of better flexibility and handling of changes. The majority of network changes take place in the edge where switches and links are added as traffic requires, whereas the core remains static. With PNNI in the edge, these changes require little or no reconfiguration of the switches involved.

Using PNNI at the core or data center and IISP at the edge or wiring closet devices has the following advantages:

- The edge or wiring closet devices require the least configuration. Route tables for these devices need only include route entries with a default route for both uplinks.
- The core switch requires a more complex configuration to learn the network topology and configure appropriate routes and alternate routes as needed. The core switches must have IISP routes with the proper configuration. PNNI then advertises this information to the PNNI cloud.

If only the data center devices utilize the PNNI protocol, call rerouting can be accomplished using IISP with PNNI. Figure 4-4 illustrates this type of network, with dual PNNI links connecting switches A and B and dual-homed IISP links providing connectivity from edge switch 1 to switches A and B. The same is true for the switch 3 connection to the data center. The IISP edge switch nodes establish the least-cost data path, and the client-server data path can use either of the direct data paths.

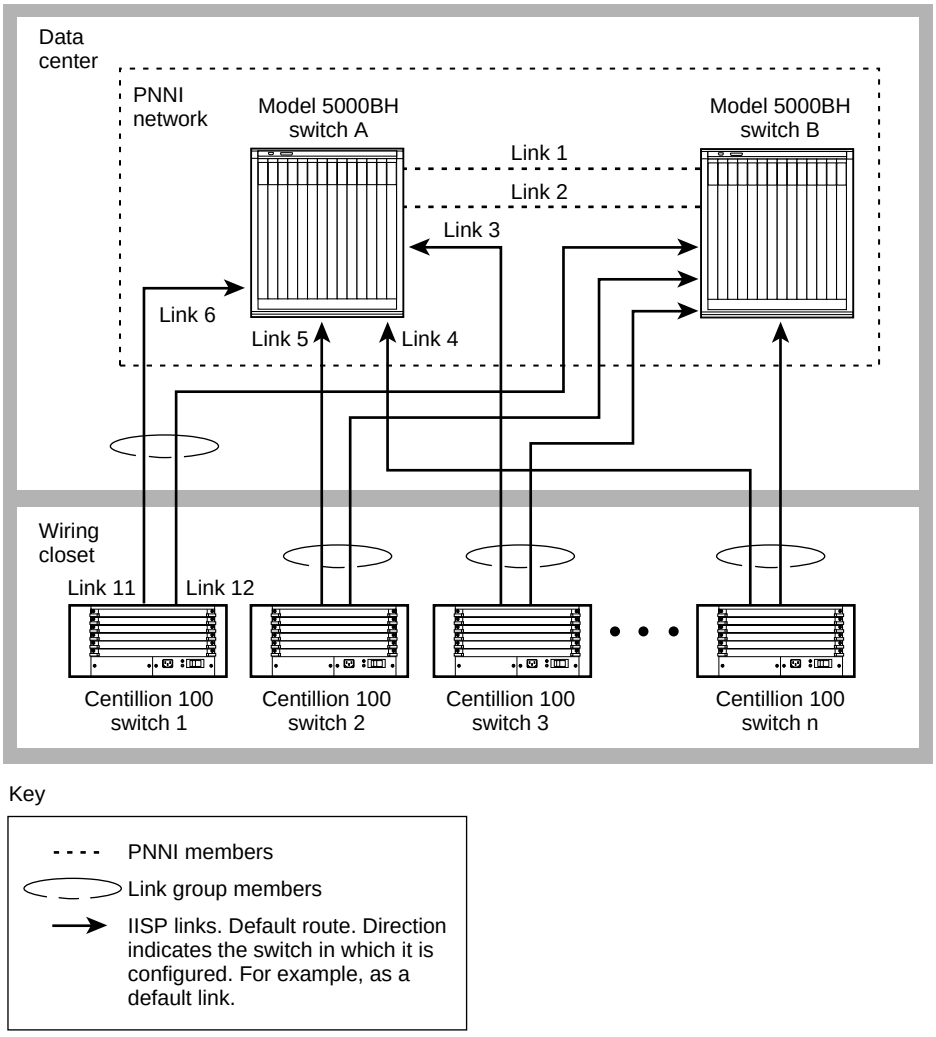

Figure 4-4. ATM Network with PNNI Core and IISP Edge

In the ATM network shown in Figure 4-4, the data center switches make up a PNNI network with IISP links to the edge switches. Table 4-2 indicates the routing information that would be required if IISP were used in the data center to define reachability to the edge switches. Tables 4-3 and 4-4 indicate the required routing information for edge switches.

Table 4-2. IISP Route Table in the Data Center Edge Switches (Switch A)

Destination Address (Switch Number)	Link (Slot:Port)	Cost
1	6	10
2	5	10
3	3	10
...		
n	4	10

Table 4-3. IISP Route Table in the Wiring Closet (Switch 1)

Destination Address	Link (Slot:Port)	Cost
Default	11	10
Default	12	10

Table 4-4. Link Group Table

Link (Slot:Port)	Cost
11	N/A
12	N/A

Problems can occur when mixing PNNI and IISP in networks if the design is not carefully thought out. Figure 4-5 contains two examples of inadequate network design.

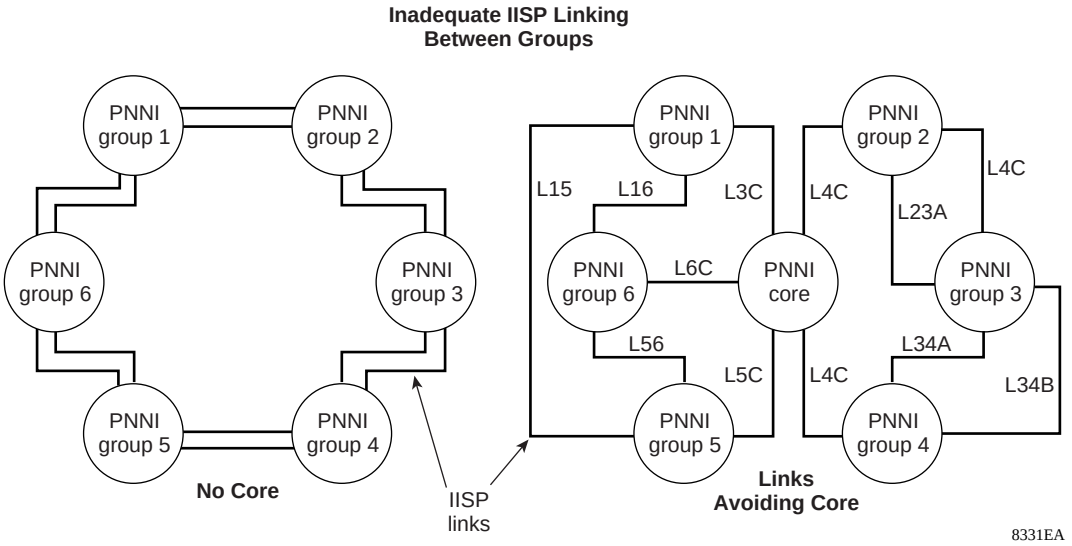

Figure 4-5. Example of Inadequate IISP/PNNI Network Designs

The No Core design in the Figure 4-5 is particularly dangerous in situations where there are more than three PNNI groups that are not interconnected in a full mesh because IISP link states will not propagate correctly between PNNI groups. For example, if the link between groups 1 and 6 fails, then groups 5 and 2 will not be aware of this and may try to connect to each other by way of groups 1 and 6, which will fail.

The Links Avoiding Core example shows two problems. The first concerns PNNI group 3, which is not connected to the core at all. PNNI group 3 has no idea of the status of the links to the core from groups 4 and 2 (L4C and L2C). For example, if L4C fails, a switch in group 3 may still attempt to route a call to the core through L34A. The call will then probably return to group 3 via L34B, creating a routing loop. If both L4C and L2C fail, all calls between any of the three groups and the core will loop continuously.

The second problem concerns PNNI groups 1, 6, and 5, which have alternate ways to connect, either by way of the core or directly to each other. For redundancy, it is likely that nonoptimum IISP routes will be specified, which may lead to an unnecessarily large routing table because of the large number of potential routes. Having large routing tables may cause problems in recovering from network failure. On the other hand, if the three groups were connected to the core twice (costing only two more ports in the core and a minor loss of potential available bandwidth), the routing table reduces to two rows, which is simpler to understand and manage.

Configuring PNNI

PNNI groups are based on switch NSAP addresses; therefore, addressing is much more important with PNNI than with IISP, especially for large networks or networks that at some future date may be connected to a public network. This fact should be taken into consideration when choosing the default level and addressing scheme. Probably the most important consideration is whether or not a registered NSAP address will be used. Not using a registered NSAP address allows more freedom, but the addressing must still be carefully planned to make routing tables more manageable.

Addressing Schemes

The full (20-byte) registered NSAP ATM address is broken down as follows:

39.CCCC.RR.AAAAAA.WW.WW.XX.XX.XX.XX.YYYYYYYYYYYY.ZZ

The various fields are defined in Table 4-5.

Table 4-5. Registered NSAP ATM Address Fields

Field	Description
39	Defines the address as a DCC ATM Format.
CCCC	Country code; for example, a value of 840F defines the address as being registered in the U.S. **NOTE:** The Data Country Code (DCC) for the U.S. is actually 840, but the specification requires this number to be left justified and padded with Fs.
RR	The registering body; for example, 80 defines this address as being registered with ANSI.

Table 4-5. Registered NSAP ATM Address Fields (continued)

Field	Description
AAAAAA	The Administrative Authority; that is, the organization that requested the registered NSAP.
WW.WW	Reserved.
XX.XX.XX.XX	Routing Domain (RD) and Area, which is the field defined to be used to address your switches.
YYYYYYYYYYYY	End station identifier (ESI).
ZZ	Selector byte.

When creating a PNNI hierarchy, therefore, the 39.CCCC.RR.AAAAAA is fixed and the user can designate the WW.WW and XX.XX.XX.XX values. For full standards compliance, the WW.WW field should be 00.00. Using a scope of 96 allows the user to configure a sample address like the one in the following example:

39.CCCC.RR.AAAAAA.00.00.DV.PL.PN.SW

Table 4-6 describes some possible user-defined fields.

Table 4-6. Sample User-Defined NSAP Fields

Field	Description
DV	A division in a conglomerate. For example, in a bank, 00 could be the holding company, 01 might be retail banking, 02 mergers and acquisitions, 03 foreign exchange, and so on.
PL	A location. For example, 00 = New York, 01 = Chicago, 02 = San Francisco. Optionally, this field could be partitioned into two parts representing city and state or city and building, depending upon the organization.
PN	The PNNI group in the building or campus, including everything that is connected via the local fiber. In a campus environment, this byte could be partitioned so that the top bits define the building and the bottom bits define the group in the building.
SW	The individual switch number in the group. Because there are 256 possible values here but only 24 to 32 switches in the group, it would be possible to partition this byte to provide meaning. For example, the top half could be the floor and the bottom half the closet so that switch 11 is in the first floor, first closet and switch 63 is in the sixth floor, third closet.

To make the hierarchy manageable, observe the following rules:

- The split between DV (division) and PL (locations) is arbitrary. Users may want split DV and PL based on organizational grounds, geographic grounds, or similar criteria. Smaller organizations will not need both. Larger organizations may want to subdivide these bytes or to extend into the reserved field.
- The split between PL and PN (PNNI group) should be the boundary of the WAN. The rule should be that if the user is charged for the use of the link between two locations, then the locations will have a different PL value.
- The core switch group in a location will be PN=00.
- In a noncore PNNI group, the switches SW=00 and SW=01 will be the transit switches to the core. Refer to Figure 4-1 on page 4-4.
- If the user has a well-defined IP addressing scheme, it may be appropriate to use that as the basis for the switch grouping, but not if IP subnets stretch across multiple locations.

Remember that if the hierarchy is planned well now, regardless of what changes may occur in the future, an organization should never have to modify the ATM address of any switch that remains in its original wiring closet and should be able to expand the number of switches in a particular closet without any complications.

Configuring PNNI with SpeedView

In setting up PNNI networks with SpeedView, you need to configure only a node ID and peer group ID for each switch. In most cases, you can use the default values and may normally change only the Node Level. For the Centillion 50/100 and Model 5000BH switches, PNNI is configured in the PNNI dialog box (Figure 4-6).

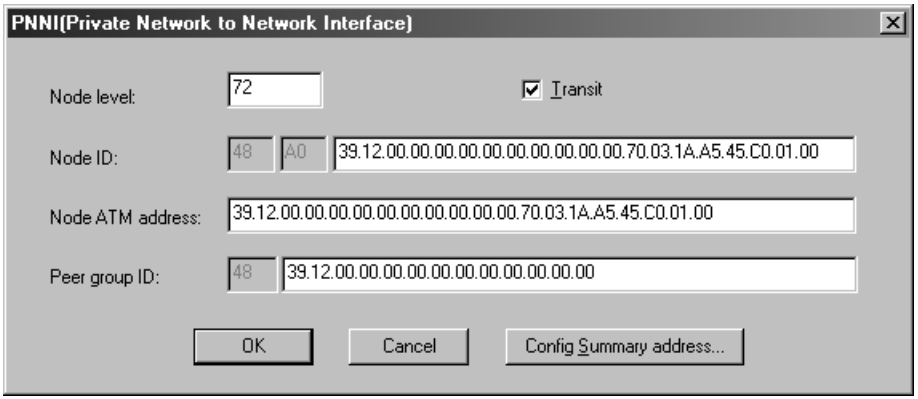

Figure 4-6. PNNI Dialog Box

Note: In most cases, default values should be used for all fields in the PNNI dialog box except the Node level. The following field explanations are for information only.

The PNNI dialog box consists of the following fields and buttons:

- Node level—the position in the PNNI hierarchy in which this node exists. This level must be consistent with the other switches in the network. The range is 1 to 104 bits with a default of 72 bits.
- Node ID—the identification number (22 bytes) of the logical node within the routing domain.
- Node ATM address—a 20-byte address. Bay Networks recommends that you use the default for this value.
- Peer group ID—the string of bits (up to 104) used to unambiguously identify a peer group.
- Transit. If you unselect this check box, a call cannot be routed through the switch.
- Config Summary address. This button allows you to configure a summary address that will be advertised instead of individual addresses. Clicking on the button opens the PNNI Summary Address dialog box, where you can enter or edit the summary address and designate interior and/or exterior address suppression.

PNNI Features

The following features are supported in the Centillion 50/100 and Model 5000BH switches:

- Logical links (this page)
- Multilevel designated transit lists (page 4-21)
- Border nodes (page 4-22)
- Address summarization (page 4-22)
- Call balancing in PNNI networks (page 4-25)

Logical Links

A link between two nodes is a logical link. A logical link may be a physical link or a virtual path or channel. PNNI supports signaling and routing over multiple virtual path connections (VPCs) to multiple destinations through a single physical interface. In this situation, each VPC configured for use as a logical link has a signaling channel associated with it. Virtual channels within these VPCs are controlled only by the associated signaling channel of that particular VPC. The default signaling channel (VPI=0) on the same physical interface does not control virtual channels within other VPCs used as logical links, but it does control all the remaining virtual channels and virtual paths on the physical link. Thus, on Centillion 50/100 and Model 5000BH switches, users can configure up to three logical links or VPCs on a given physical interface. This feature is helpful when users want to interconnect geographically separated locations through a WAN, as shown in Figure 4-7.

Caution: You cannot configure an ELAN (LECs and LES/BUS) that spans multiple logical links on the same physical interface. For example, in Figure 4-7, you cannot have LEC (switch B) and LEC (switch C) on the same ELAN because the switch is unable to set up a point-to-multipoint (PMP) SVC that branches on more than one VPI on the same port. If LECs on the two remote switches attempt to join the same ELAN, then only the first one will be successful because it is not possible to set up the relevant add party requests for the second LEC.

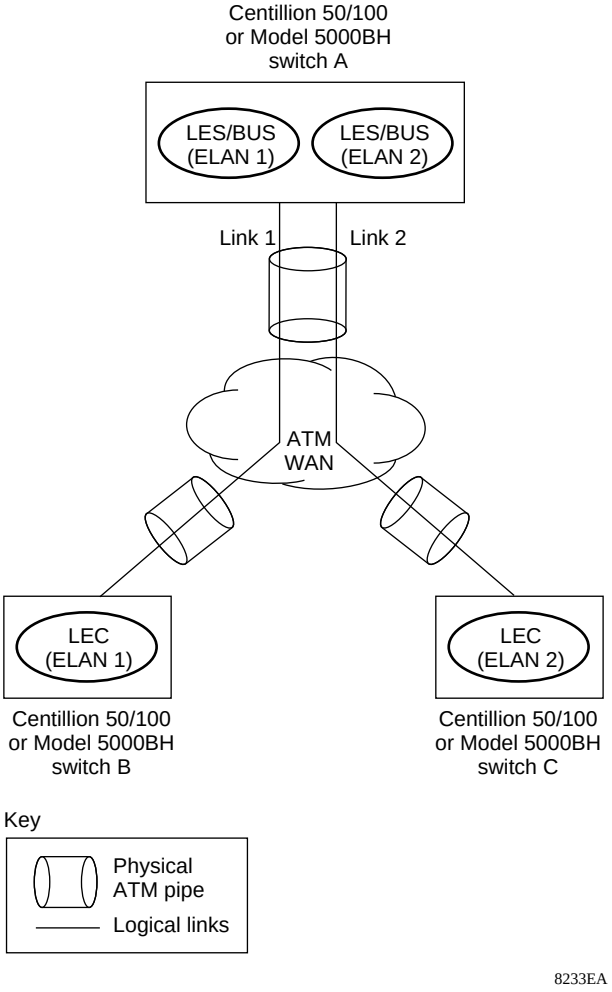

Figure 4-7. Logical vs. Physical Links

Multilevel Designated Transit Lists

The Centillion 50/100 switches support multilevel designated transit lists. Multilevel DTLs include peer groups in different levels of the PNNI hierarchy, which means that end stations connected to Centillion 50/100 switches in separate peer groups can communicate over PNNI.

Figure 4-8 shows three peer groups (A, B, and C). End station Z can talk to end station Y because switch 2 in peer group A and switch 3 in peer group B create a third peer group (C) that allows peer groups A and B to communicate.

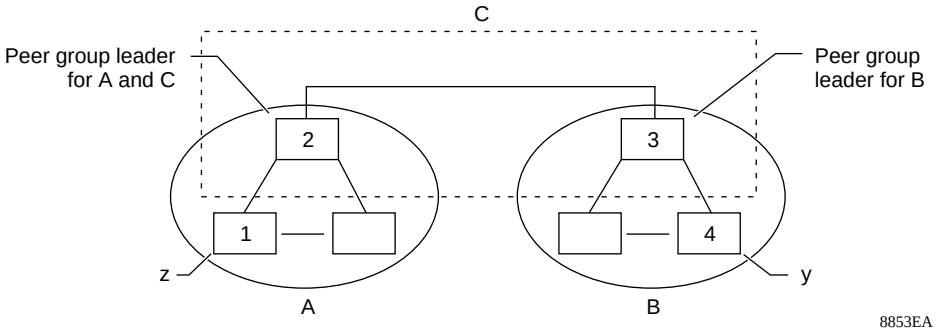

Figure 4-8. Multilevel Hierarchy of Centillion 50/100 Switches

Switch 2 is the peer group leader of peer group A and C. Switch 2 cannot be the peer group leader of B because a switch cannot be the peer group leader of another peer group on the same level.

Communication between end stations in different peer groups is accomplished using a multilevel DTL, which includes nodes in different levels of the PNNI hierarchy. The DTL of the path from station Z to station Y would be: 1-2-C-3-4.

The DTL is considered multilevel because it includes peer group C (a higher-level peer group).

Note: A Centillion 50/100 or Model 5000BH switch cannot be a peer group leader or border node.

Border Nodes

As explained previously, a collection of nodes (switches) for a peer group appears from outside of the peer group to be a single node, referred to as a *logical node.* Within a peer group, all nodes are peers and see the full topology of the peer group. Neighbors that discover they belong to different peer groups do not exchange topology information and automatically become *border nodes*. A border node can be a logical node or a physical device in a specified peer group that has at least one link that crosses the peer group boundary. Figure 4-9 illustrates a peer group with five border nodes (the black dots). Links that cross peer groups are called *outside links* that continue to exchange hellos.

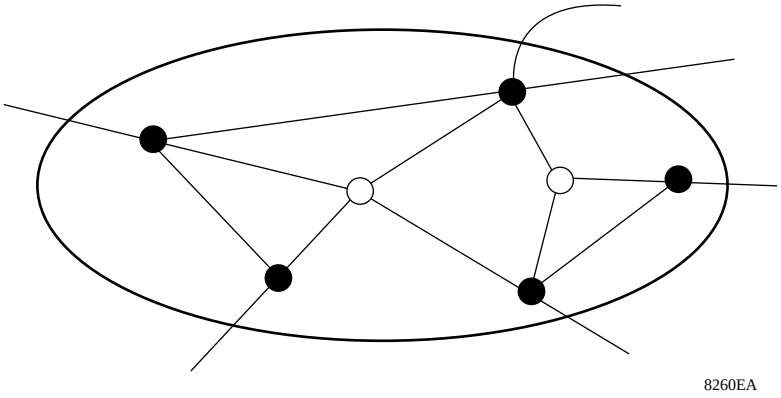

Figure 4-9. Peer Group with Five Border Nodes

Address Summarization

Generally, PNNI advertises end system reachability using prefixes on ATM addresses to form summaries, making it possible to advertise reachability to a large number of end systems using a single prefix. If an end system attached to a node does not fit into one of the node's configured summaries, it will be necessary for the node to make an explicit advertisement for that end system. Using address summarization makes the PNNI network more scalable by reducing the number of addresses it must advertise. By properly configuring the internal summary addresses, you can keep the amount of reachability information low.

Figure 4-10 shows how address summarization works. The switch prefixes and link numbers are shown for area A. Assume that default summarization is enabled in all the switches. Switches A1 and A2 can learn reachability information from switches A01 through A120 by using either static configuration or PNNI. If these switches are not running PNNI, users can configure link 1 A01 through link 1 A120 and link 2 A01 through link 2 A120 as IISP links and then statically configure the routes. In switches A1 and A2, specific routes for the addresses 39.00.01.0A.00.01 through 39.00.01.0A.00.20 are configured. These addresses are seen as external addresses (exterior reachable addresses).

If switches A01 through A120 are running PNNI, then switches A1 and A2 will learn the addresses of switches A01 through A120 by exchanging node IDs. In this case, the addresses 39.00.01.0A.00.01 through 39.00.01.0A.00.20 are seen as interior reachable addresses. If switches A1 and A2 have only default summarization turned on, they must advertise all these addresses to the other switches (B1, B2, C1, C2) in the PNNI group. Thus the number of reachable addresses in the network is very high whether switches A01 through A120 are running PNNI or IISP.

PNNI address summarization can cut down the number of reachable addresses advertised. If a summary address (interior or exterior) 39.00.01.0A is configured in switches A1 and A2, then these switches will advertise only this address as it summarizes the specific addresses 39.00.01.0A.00.01 through 39.00.01.0A.00.20. Thus the number of addresses advertised to switches B1, B2, C1, and C2 is reduced from 20 to 1.

By properly configuring the internal summary addresses, you can run PNNI on edge switches also but keep the advertised reachability information to a low level. For example, you can configure an internal summary address 39.00.01.0A in switches A1 and A2, 39.00.01.0B in switches B1 and B2, and 39.00.01.0C in switches C1 and C2.

Configured summaries provide the capability either to advertise or suppress matching addresses. Therefore, a user wanting to screen an address for security reasons may configure a matching summary with suppression enabled.

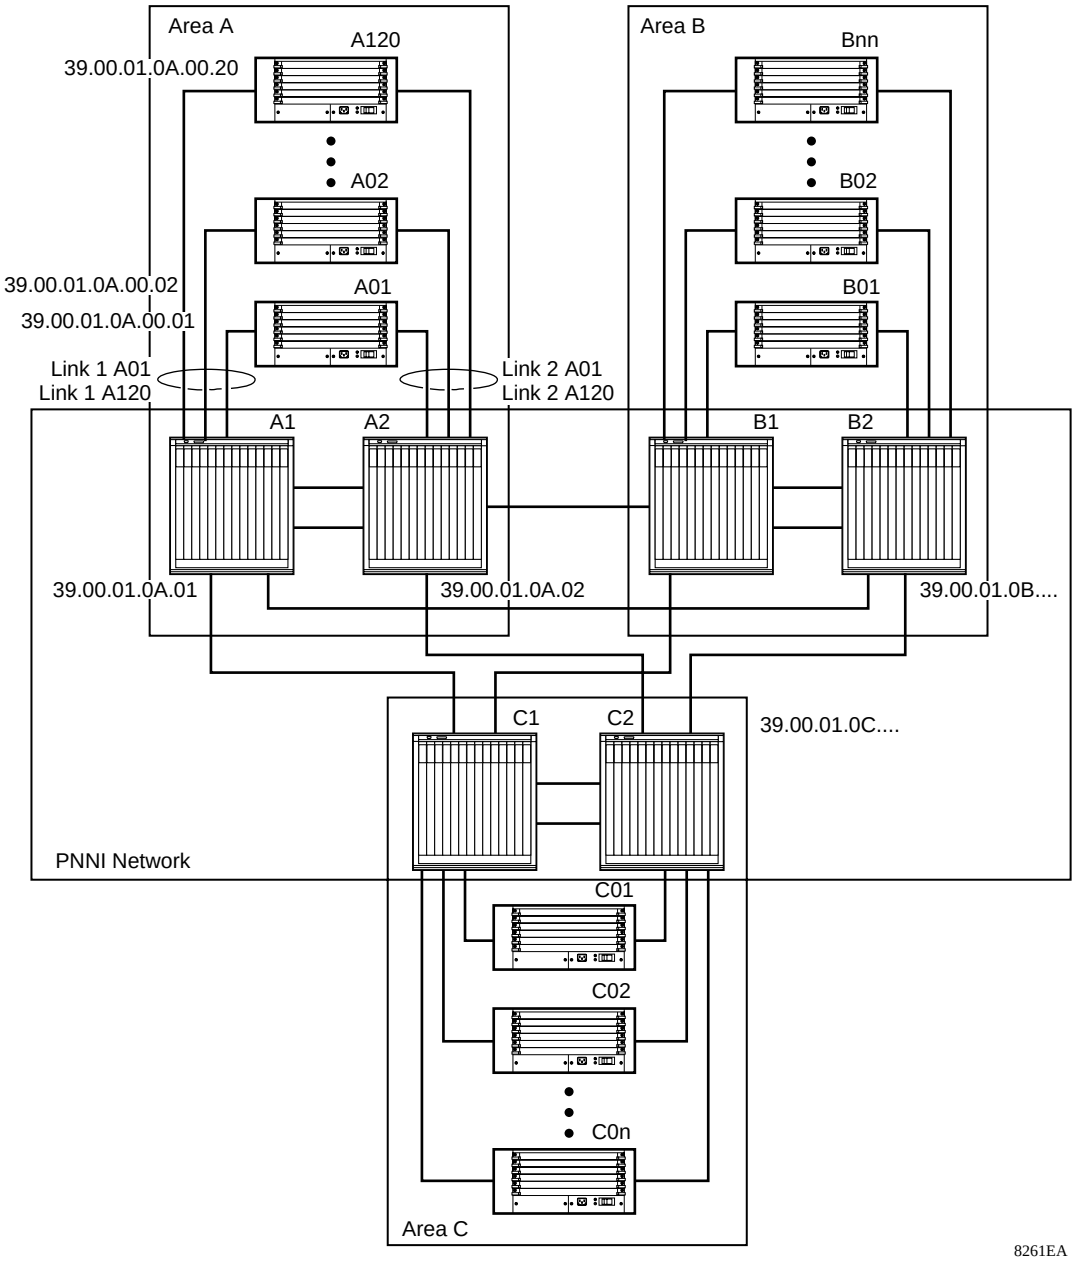

Figure 4-10. Address Summarization

Call Balancing in PNNI Networks

In the current version of advanced switch image, PNNI supports call balancing across equal cost paths and across equal cost hops if they lie within a path selected by the originating node. The call balancing algorithm employed by PNNI attempts to distribute the calls evenly across equal cost paths and uses a randomization algorithm to distribute the calls across equal cost hops lying within a selected path.

Figure 4-11 shows an example where calls can be routed from switch 1 to switch 4 along two different equal cost paths: switches 1-2-3-4 or switches 1-5-6-4.

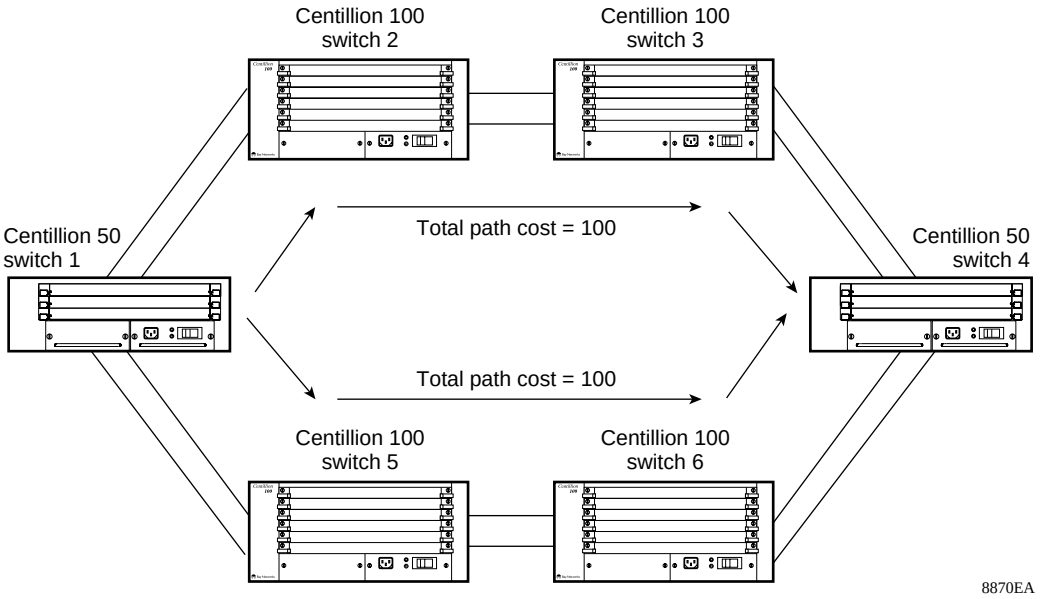

Figure 4-11. PNNI Call Balancing

Setting up 100 calls will result in 50 calls taking the first path and the remaining 50 taking the second path. The sum total of all the calls originating out of switch 1 toward switch 5 is 50 as are the total number of calls originating out of switch 1 toward switch 2.

The assumption is that redundant links along each hop of a selected path have the same cost. The total calls along each redundant link on a given hop are randomly distributed. That is, along the path with switches 1-5-6-4, if the two links between switch 1 and switch 5 have the same cost associated with them, then the calls are randomly distributed between the two links—however, the total number of calls should equal 50.

Chapter 5
IISP in an ATM Network

This chapter describes the Bay Networks implementation of the ATM Forum Interim Inter-switch Signaling Protocol (IISP) for the private network-to-network interface (PNNI) in the ATM network. IISP is based on the UNI 3.1 signaling protocol and provides optional support for the UNI 3.0 signaling protocol. SpeedView and the Centillion 50/100 and Model 5000BH switches support the UNI 3.0 and 3.1 signaling protocols and provide full support of the IISP specification. For information about PNNI and using PNNI along with IISP, refer to Chapter 4, "PNNI in an ATM Network."

This chapter discusses the following topics:

- Where IISP fits in an ATM network (this page)
- IISP link configuration (page 5-3)
- Link groups and loop prevention (page 5-6)

Where IISP Fits in an ATM Network

IISP is the first protocol that enables SVC connectivity between multivendor ATM switches. Prior to IISP, the only standard way to communicate between ATM switches was over PVCs and PVPs. PVCs required every circuit between end users to be set up manually. With IISP, only the switch-to-switch routing tables need manual configuration.

IISP is in the Centillion 50/100 and Model 5000BH switch images. SpeedView can be used to configure IISP within each switch. The Centillion 100 and Model 5000BH switches can be configured as the network side or the user side of an IISP link.

Figure 5-1 illustrates where IISP fits in an ATM network.

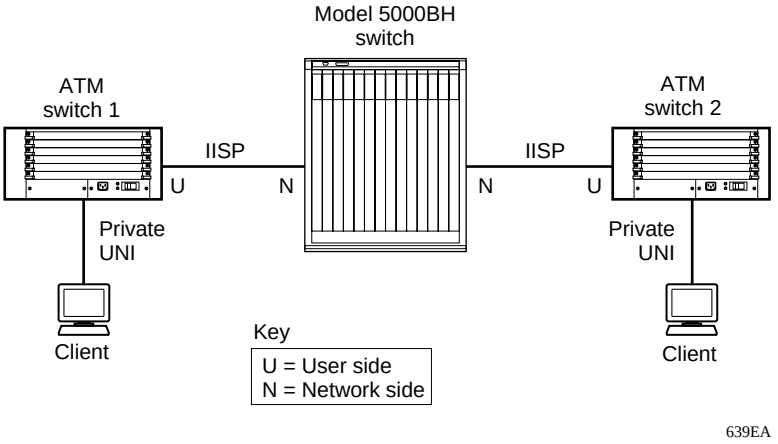

Figure 5-1. IISP Between Private ATM Network Switches

Call collisions on an IISP link occur if two ATM switches assign identical VPI/VCIs to calls simultaneously. To prevent this problem, the network side is given the responsibility of assigning VPI/VCIs for calls.

IISP Link Configuration

Routing is accomplished on a hop-by-hop basis by making a best match of the destination ATM address in the call setup request with ATM address entries in the IISP routing table at each switch. Before IISP can be used for call processing in an ATM network, a static routing table must be set up for each Centillion 50/100 or Model 5000BH switch. A static routing table is required because routing information is not exchanged dynamically between switches over an IISP link.

For the Centillion 50/100 and Model 5000BH switches, the routing table is based on the routing information you enter in the Configure Call Route dialog box. You must provide and manually configure the IISP routing entries for each destination ATM address. Each IISP port on an ATM switch requires an IISP route entry.

There can be one or multiple IISP route entries for each port that carries traffic. Also, each Centillion 50/100 and Model 5000BH ATM switch needs a separate routing table if IISP is going to be used.

IISP Route Table

The IISP Route Entry Table (Figure 5-2) defines IISP links for Centillion 50/100 and Model 5000BH switches.

The IISP Route Entry Table consists of the following parameters:

- ATM Address—the destination ATM address for the IISP link
- Route Type—the type of the link
- Module—the switch module for the IISP link
- Port—the switch module's port to which the IISP link is connected
- VPI—the virtual path identifier (VPI) of the link
- Cost—the cost for the link

  The cost is an administrative weight placed on a link to determine the order in which it will be used. A link with a lower cost has a higher priority and will be used before a link with a higher cost.

- Link Group number (optional)—an integer that is assigned for each respective link group

  Links assigned the same number are part of the same link group. In the call routing decision, the calls will not be routed over links that are in the same link group. Not routing a call over links that are in the same link group helps eliminate potential loops in the network. Assigning two ports within a switch to the same link group causes the two ports to behave as a single physical link. For example, no data that is received on a port within a link group will ever be transmitted out another physical port that is a member of that same link group.

- Scope (not used; PNNI only)

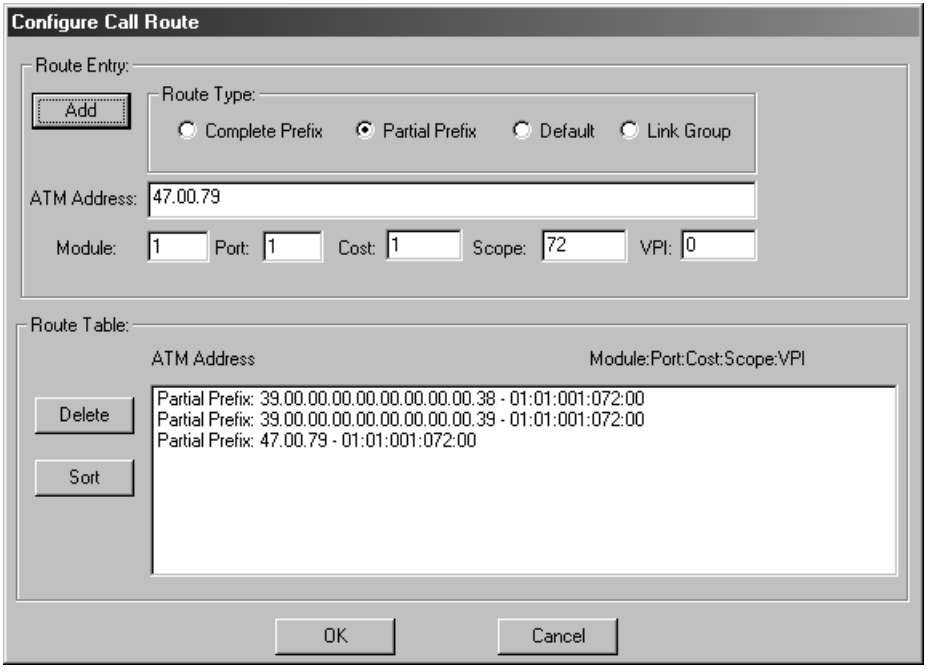

Figure 5-2. Configure Call Route Dialog Box with Route Entry Table

Route Entry Table

The routing entry information consists of the following four parameters:

- ATM Address
- Route Type (Complete Prefix, Partial Prefix, Default, Link Group)
- Module and Port
- Link Cost

ATM Address

The ATM address can be either a complete 19-byte network service access point (NSAP) address or an ATM prefix.

Route Type

The Complete Prefix radio button is selected when the entire 19-byte NSAP host address is used to route a call directly to that host.

The Partial Prefix radio button is used when only the ATM prefix, part of the ATM prefix, or an incomplete NSAP is used to route calls to a specific destination. Any calls destined for a host with this prefix will be routed to that network.

The address length or mask is determined by the choice of either Complete or Partial Prefix and is automatically set to the location of the last nonzero bit in the address.

Module and Port

These entries identify the module (slot) and the port for the IISP link. The module and port identify the egress port to which the call should be routed.

The entries consist of the following values:

- Module number—the switch module's number for the IISP link
- Port number—the switch module's port number for the IISP link

Link Cost

The link cost is an optional user-assigned value that can be used to assign link priority. Depending on the network configuration, it is possible for the link selection process to discover two or more links that have equal lengths. If IISP discovers multiple links of equal length during a call setup procedure, IISP uses the link cost value to select one link. The call is routed over the link that has the lowest cost. If the cost is equal for the links, calls are routed over the links in a randomized round-robin fashion so that calls are balanced between the links.

Link Groups and Loop Prevention

Alternate and parallel routes offer protection against link or switch failure and allow you to prioritize links. However, loops can occur in networks that use alternate and parallel links. The Centillion 50/100 and Model 5000BH switch software provides a functionality called “link groups” to prevent the formation of loops in the network. A link group is an association of specific IISP links. The IISP links are manually assigned to a user-defined link group.

To prevent a call request from continually being forwarded in a loop, the Centillion 50/100 and Model 5000BH switch software does not allow a call received on an IISP link that is assigned to a specific link group to be routed over another IISP link assigned to the same link group. For example, if a call is received on an IISP link assigned to link group 1, the IISP procedure will not route the call on another IISP link that belongs to link group 1. In the implementation of IISP, this technique prevents loops.

An IISP link is assigned to a link group by selecting the link group radio button in the Route Type field of the Configure Call Route dialog box (Figure 5-2) when you configure your IISP Route Entry Table.

How a Link Group Provides Loop Prevention

The ATM network topology in Figure 5-3 consists of alternate and parallel IISP links. The assumption for this scenario is that client A is setting up a connection to the server and switch 3 is down before the call request from client A is initiated.

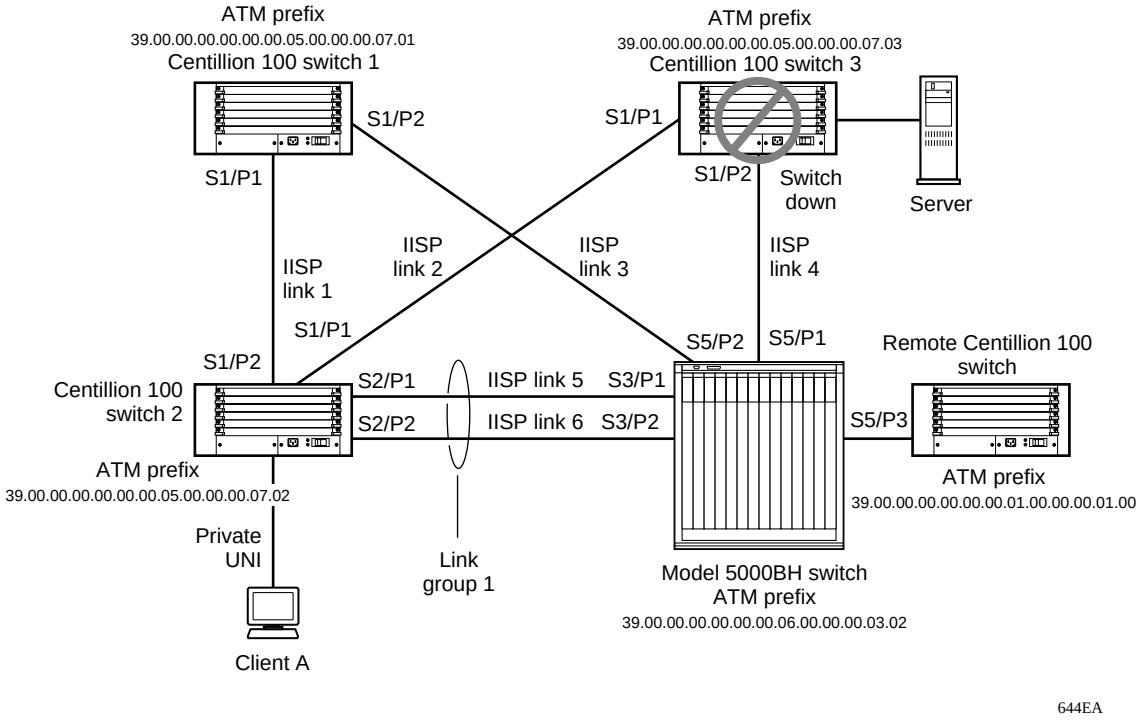

Figure 5-3. Example of Link Groups and Loop Prevention

When client A places a call request to the server, the IISP procedure checks for link availability and determines that IISP link 2 is down. The IISP procedure searches the routing table (Table 5-1) for an alternate route by selecting the next best match to the ATM address of the server.

Table 5-1. Centillion 100 Switch 2 Routing Table

Destination ATM Address	Cost	Interface Index	
		Slot	Port
39.00.00.00.00.00.00.05.00.00.00.07.01	3	1	2
39.00.00.00.00.00.00.05.00.00.00.07.03	2	1	1
39.00.00.00.00.00.00.06	4	2	1
39.00.00.00.00.00.00.06	4	2	2
39.00.00.00	4	2	1
39.00.00.00	4	2	2
00.00.00...		2	2

Also notice that Centillion 100 switch 2 has IISP links 5 and 6 (Table 5-2 and Figure 5-3) assigned to link group 1.

Table 5-2. Switch 2 Link Groups

Slot	Port	Link Group
2	1	1
2	2	1

Multiple addresses can match the call setup destination address. In this case, the following four routing table entries that match are listed in order of preference with the first entry being the most preferred, the second and third being equal, and the fourth being the route of last resort:

- 39.00.00.00.00.00.00.05.00.00.00.07.03 (but the link is down because Centillion 100 switch 3 is down)
- 39.00.00.00 (slot 2, port 1, of Centillion 100 switch 1) for IISP link 5
- 39.00.00.00 (slot 2, port 2 of Centillion 100 switch 2) for IISP link 6
- 00.00.00... (default route of last resort)

Because the length and cost are equal for IISP links 5 and 6, the IISP procedure selects the link with the fewest calls. For this example, assume that IISP link 5 is selected as the route to the next hop toward the final destination.

When the call reaches the Model 5000BH switch in Figure 5-3, the IISP procedure searches the routing table (Table 5-3) to determine the next route to the destination.

Table 5-3. Model 5000BH Switch Routing Table

Destination ATM Address	Cost	Interface Index	
		Slot	Port
39.00.00.00.00.00.00.05.00.00.00.07.01	1	5	2
39.00.00.00.00.00.00.05.00.00.00.07.02	1	3	1
39.00.00.00.00.00.00.05.00.00.00.07.02	1	3	2
39.00.00.00.00.00.00.05.00.00.00.07.03	1	5	1
39.00.00.00.00.00.00.05	4	3	1
39.00.00.00.00.00.00.05	4	3	2
39.00.00.00.00.00.00.01	4	5	3

The IISP link for Model 5000BH switch slot 5, port 1 (link 4) is down; therefore, the IISP procedure searches the Model 5000BH switch routing table for an alternate route.

At first glance at the routing table, three alternate routes to the final destination appear to be available:

- 39.00.00.00.00.00.00.05.00.00.00.07.03 (but the link is down because the Centillion 100 switch 3 is down)
- 39.00.00.00.00.00.00.05, slot 3, port 1
- 39.00.00.00.00.00.00.05, slot 3, port 2

However, the IISP procedure also checks the IISP interface link table (Table 5-4) for link group assignments.

Table 5-4. IISP Link Groups

Slot	Port	Link Group
3	1	1
3	2	1

The call request was received on slot 3, port 1 (IISP link 5), which is assigned to link group 1; therefore, the IISP procedure does not forward the call over IISP link 6 or back on link 5 because link 5 and link 6 are assigned to link group 1. In this case, the call request is dropped as it should be, and the system returns a message of “no route to destination.”

Notice that the Model 5000BH switch routing table does not contain an entry for a route to switch 3 through switch 1. The reason is that if such a route did exist, the call might be circulated in an endless loop. One possible loop path might exist solely over IISP link 3. Another possible loop might develop between Model 5000BH switch 1 and switch 2.

You can avoid causing loops by simply not providing routes. But if you need to design your network to have redundant routes, check that there are no potential loops. Also, minimize the assignment of links to link groups. Using a link group assignment incorrectly can result in a call request not being forwarded.

Example of IISP and Redundant Links in a Network

The ATM network topology shown in Figure 5-4 provides an example of redundant links and dual-homed network connectivity.

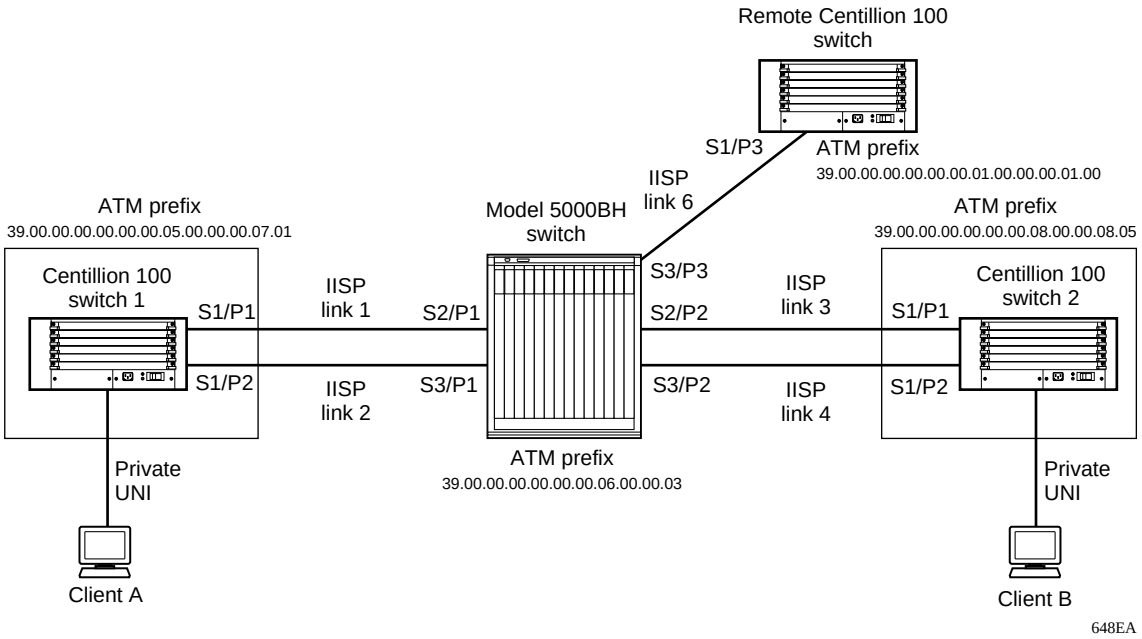

Figure 5-4. Example of IISP Connectivity and Redundant Links

The routing table entries for switch 1 are shown in Table 5-5. The redundant IISP links are between switch 1 and the Model 5000BH switch and between switch 2 and the Model 5000BH switch. Both addresses have the same length match, but IISP link 1 and link 2 have different link costs. In this case, IISP link 2 is the preferred link, because the link cost is 2. IISP link 1 is the redundant link.

Table 5-5. Centillion 100 Switch 1 Routing Table

Destination ATM Address	Cost	Interface Index	
		Slot	Port
Default	3	1	1
Default	2	1	2

When client A places a call to client B, the IISP procedure routes the call request over IISP link 2. If IISP link 2 becomes unavailable, the IISP procedure routes the call request over link 1.

If IISP links 1 and 2 were assigned the same cost, the IISP procedure would perform call balancing across the links.

The routing table for the Model 5000BH switch in Figure 5-4 is shown in Table 5-6.

Table 5-6. Model 5000BH Switch Routing Table

Destination ATM Address	Cost	Interface Index	
		Slot	Port
39.00.00.00.00.00.00.05.00.00.00.07	2	2	1
39.00.00.00.00.00.00.05.00.00.00.07	3	3	1
39.00.00.00.00.00.00.08.00.00.00.08	3	2	2
39.00.00.00.00.00.00.08.00.00.00.08	2	3	2
39.00.00.00.00.00.00.01.00.00.00.01	3	3	2

The routing table for Centillion 100 switch 2 in Figure 5-4 is shown in Table 5-7.

Table 5-7. Centillion 100 Switch 2 Routing Table

Destination ATM Address	Cost	Interface Index	
		Slot	Port
Default	3	1	1
Default	2	1	2

Example of Link Groups in a Network

Figure 5-5 shows a network of Centillion 50/100 and Model 5000BH switches with IISP links assigned to link groups.

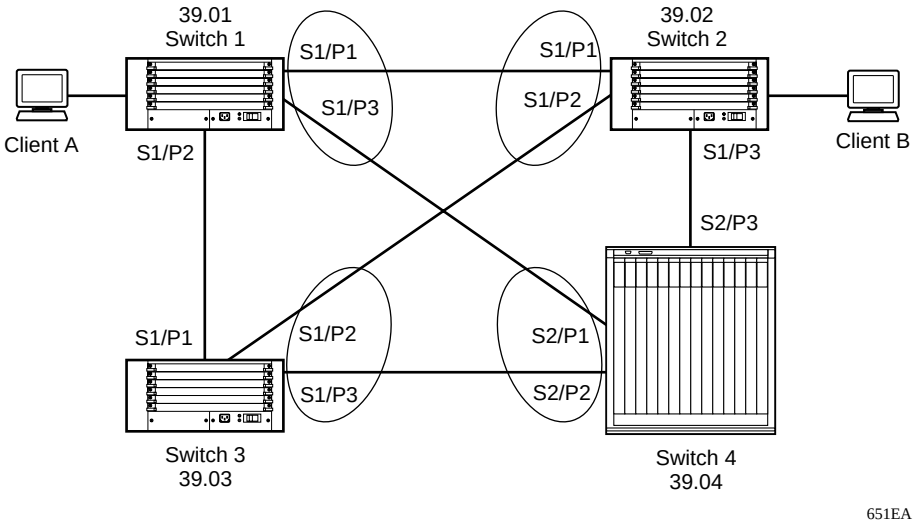

Figure 5-5. Link Groups in an IISP Network

Each switch shown in Figure 5-5 has its own routing table. Each switch also has link groups configured for the different modules and ports.

The switch routing tables and link groups for each switch in Figure 5-5 are listed in Table 5-8 through Table 5-15.

Table 5-8. Switch 1 (39.01) IISP Routing Table

Destination ATM Address	Cost	Interface Index	
		Slot	Port
39.02	4	1	1
39.03	4	1	2
39.04	4	1	3
00.00		1	2

Table 5-9. Switch 1 (39.01) Link Groups

Slot	Port	Link Group
1	1	1
1	3	1

Table 5-10. Switch 2 (39.02) IISP Routing Table

Destination ATM Address	Cost	Interface Index	
		Slot	Port
39.01	4	1	1
39.03	4	1	2
39.04	4	1	3
00.00		1	3

Table 5-11. Switch 2 (39.02) Link Groups

Slot	Port	Link Group
1	1	1
1	2	1

Table 5-12. Switch 3 (39.03) Routing Table

Destination ATM Address	Cost	Interface Index	
		Slot	Port
39.01	4	1	1
39.02	4	1	2
39.04	4	1	3
00.00		1	3

Table 5-13. Switch 3 (39.03) Link Groups

Slot	Port	Link Group
1	2	1
1	3	1

Table 5-14. Switch 4 (39.04) IISP Route Table

Destination ATM Address	Cost	Interface Index	
		Slot	Port
39.01	4	2	1
39.02	4	2	3
39.03	4	2	2
00.00		2	1

Table 5-15. Switch 4 (39.04) Link Groups

Slot	Port	Link Group
2	1	1
2	2	1

In normal operation, if client A tries to set up a call for client B, the call is sent directly to switch 2 from switch 1 using the IISP route tables shown in Table 5-8 through Table 5-15.

If the link between switch 1 and switch 2 goes down (Figure 5-6), an alternate route must be used.

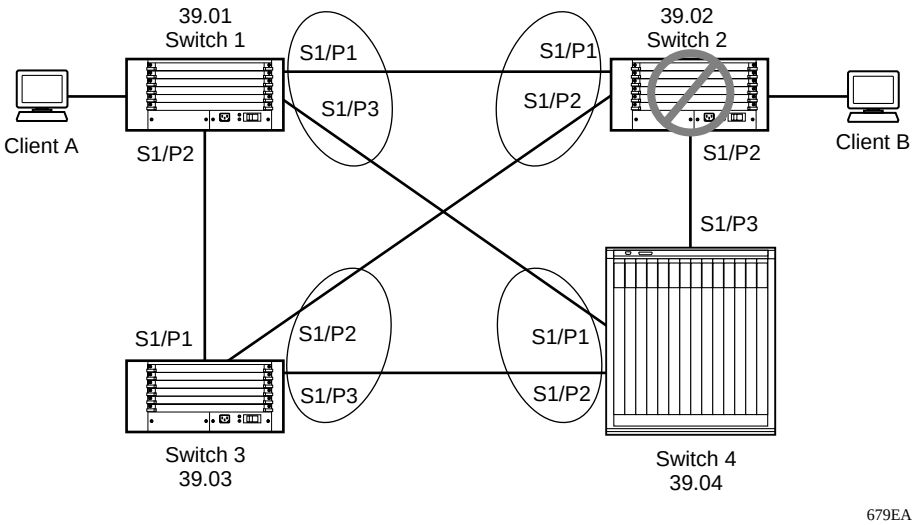

Figure 5-6. Break in the IISP Network

If client A tries to set up a call to client B, but the direct link from switch 1 to switch 2 is down, the following actions are taken to set up and deliver the call:

- Switch 1 uses its default route to client B and goes to switch 3.
- Switch 3 receives the call request and tries to send the call to switch 2. Switch 3 discovers that the direct link to switch 2 is down.
- Switch 3 uses its default route to switch 4.
- Switch 4 receives the call request and tries to send the call to switch 2. Switch 4 discovers that the direct link to switch 2 is down.
- Switch 4 would send the call request to switch 1, but S1/P2 and S1/P1 are in the same link group. Call requests that come in one interface cannot be sent out another interface that is part of the same link group. Switch 4 breaks the IISP call routing loop by not sending the call request to switch 1.

Chapter 6
Packet Filter Applications

The Centillion 50/100 and Model 5000BH switches provide powerful packet filtering. They can recognize any pattern within the first 255 bytes of a frame and take the action you specify for frames that match or do not match the pattern. Ports accept up to 64 filters (32 inbound and 32 outbound). The switch offers a NetBIOS filter template that lets you filter NetBIOS names.

You can set up filters for frames entering or leaving a port. Frames that meet the specified condition may be forwarded as usual, forwarded to additional destinations, forwarded to a monitoring port, dropped, or subjected to another filter. You can define a separate set of actions for frames that do not meet the specified condition.

Centillion packet filtering offers the following advantages:

- Reducing traffic across the network
- Analyzing network performance by routing a copy of the real-time traffic to a LAN analyzer attached to an Ethernet port
- Monitoring user backbone utilization by copying traffic to a usage accounting server in addition to the intended destination
- Controlling traffic flow by redirecting frames that meet certain criteria to a specified destination instead of or in addition to the destination addresses

For information about setting up packet filters with SpeedView, refer to *Using SpeedView 4.0 for Windows*. For more information about the NetBIOS Protocol, see Chapter 7, “NetBIOS Protocol.”

This chapter contains examples of how packet filtering can be applied to control traffic on the Centillion 50/100 and Model 5000BH switches. Use the examples in this appendix as guides to understand packet filtering, not as concrete examples applicable to all network environments. To ensure filter effectiveness, first take a trace on the LAN where you want to apply filters, and then develop them according to your specific traffic pattern.

This chapter discusses the following topics:

- Traffic management in switched internetworks (this page)
- Introduction to packet filtering (page 6-3)
- Packet filter examples (page 6-6)
- Common SAPs (page 6-27)
- Ethernet protocol types (page 6-28)

Traffic Management in Switched Internetworks

As LAN switching is deployed in collapsed backbones and microsegmented workgroups, the traffic flow within the network changes. Servers are on different segments from their clients, and they may even be on different switches. Most likely they are in a server farm somewhere in the network.

This topology change has a significant impact on the volume of broadcast and explorer traffic generated and then propagated through the network as each client tries to locate a server. Chatty protocols and uncontrolled broadcast storms eat up the bandwidth of your backbone, especially in peer-to-peer environments.

NetWare servers periodically broadcast the services they are running. In a large multiserver NetWare environment, this process can have a significant impact on available bandwidth for carrying user data. When multicast video streams, video conferencing, and other time-sensitive traffic is flowing through the network, traffic management will become all the more crucial for preserving backbone bandwidth and maximizing its payload.

The Centillion 50/100 and Model 5000BH switches are designed for large multiprotocol internetworks. These switches offer a comprehensive set of traffic management capabilities, including broadcast explorer suppression, a source route broadcast explorer proxy, a NetBIOS name proxy, and a NetBIOS name filter. All operate in conjunction with switch-wide virtual ring and network-wide VLAN capabilities, as well as network-wide VLANs under ATM LANE.

Introduction to Packet Filtering

Centillion packet filtering is designed to allow a high level of flexibility to give users better control when directing frames. Most production networks have broadcast-intensive protocols that consume bandwidth; however, with packet filtering, SpeedView gives network managers control over how broadcast frames are forwarded. Applying packet filtering can prevent wasted bandwidth and boost overall network throughput.

Packet filters on the Centillion 50/100 and Model 5000BH switches enable you to offset a specified number of bytes into a frame to locate the bytes to match. Actions are then carried out based on the matching conditions.

Token ring frames (Figure 6-1) have two offsets. The first is a MAC offset that starts at the access control (AC) field of a token ring frame. The second is an LLC offset that starts at the destination service access point (DSAP) field.

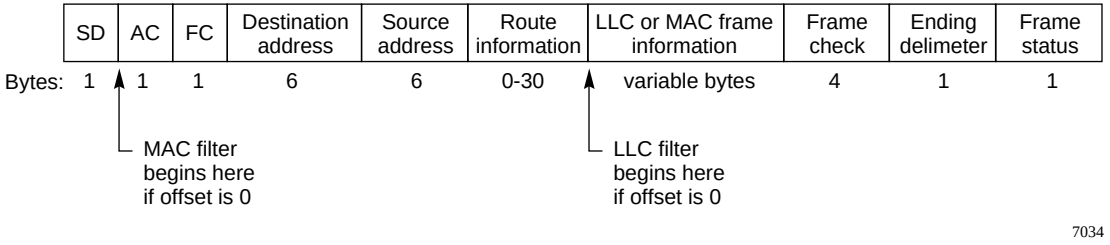

Figure 6-1. Token Ring Packet Filter Offsets

An expanded view of the LLC frame is shown in Figure 6-2.

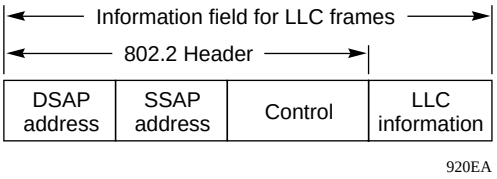

Figure 6-2. Information Field for an LLC Frame

For Ethernet frames (Figure 6-3), the only type of offset is a MAC offset, and it starts at the DA field.

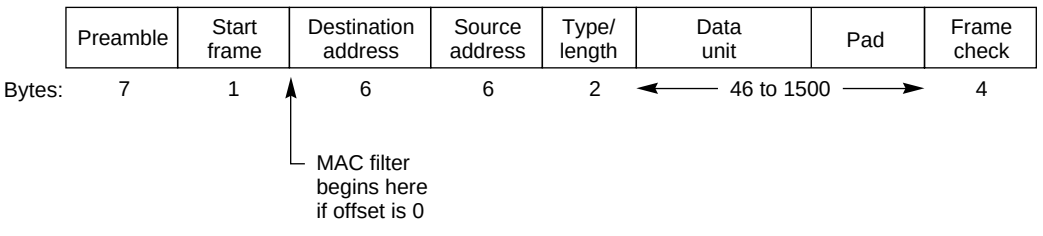

Figure 6-3. Ethernet Frame Format Showing Start of Filter Offsets

Depending on the application, Ethernet stations may operate in either Ethernet V2.0 or Ethernet 802.3 mode. Ethernet V2.0 frames use the Type field to indicate protocols used, and Ethernet 802.3 frames use the length field to indicate the frame byte count.

A filter consists of sequences that specify conditions to be checked and actions to be carried out based on matching or failing the conditions. A filter name must be assigned to each group of sequences. Upon matching a condition, the user may specify an action; if the match fails, however, the user may specify a different action.

Packet filtering on the Centillion 50/100 and Model 5000BH switches is executed on inbound or outbound traffic coming from the network. A filter can read up to 255 bytes deep into a frame and match on any consecutive 12 bytes. Sixteen sequences can be written in one filter group, and approximately 64 filters can be applied per port. Filters are processed sequentially according to the order in which they are applied on each interface.

Figure 6-4 uses a token ring filter with both MAC and LLC offsets to demonstrate the logic behind Centillion filtering. An Ethernet filter uses the same logic, but it uses only MAC offsets. Filter FL_TR (Figure 6-4) drops any token ring frame with destination address of all Fs with DSAP not equal to E0 coming into any port with FL_TR applied. Other frames are forwarded as normal.

Name	ID	Seq	Type	Offset	Value (hex)	Cond	Match	Fail	Fwd	Mon Dest	Addl Dests
FL_TR	1	1	MAC	2	FF.FF.FF.FF.FF.FF	EQ	2	0	DROP		
		2	LLC	0	E0	EQ	0	255	NORM		

Figure 6-4. Filter That Drops All Fs Broadcast with DSAP and SSAP of E0

The token ring filter FL_TR, has two sequences that check for two criteria.

Before a filter is carried out, the initial forwarding state is always set to NORM or normal (fwd=NORM). As matching conditions are met, the Forwarding state may be reinitialized according to the specified forwarding action.

Sequence 1 in FL_TR is a MAC offset of 2 bytes from the AC field that points to the destination address in a token ring frame. If the value to match (for example, FFFFFFFFFFFF) is met, the frame goes to sequence 2 of the filter group and the frame Forwarding state is reinitialized to DROP (fwd=DROP). If sequence 1 fails, the frame is forwarded and does not continue to the next applied filter on the port (0 =stop). The Forwarding state stays at NORM.

Sequence 2 is an LLC offset of 0 bytes and is carried out only if the matching condition in sequence 1 is successful. If the value to match is equal to E0, the frame stops. It does not continue to the next applied filter group because the matching instruction is a 0. Because there is a match in Figure 6-4, the Forwarding state is reinitialized again to NORM.

If the value were not equal to E0, the frame would continue to the next applied filter group. (A number of 255 or any value greater than the last sequence number of the filter group instructs the frame to go to the next applied filter group on the port.) The Forwarding state remains at DROP and is not reinitialized when a condition fails.

If another filter is applied on this port, the initial frame Forwarding state is carried over from the first filter. If the last condition fails to match, the DROP forwarding state is carried over. If there is no other applied filter on the port, the fail condition performs the last frame Forwarding state, which is to drop the frame.

The following sections provide more examples of how to filter a particular type of traffic. Use these examples with appropriate modifications if you need to apply them to your switch.

Packet Filter Examples

This section provides examples of different filter applications. The following examples are provided:

- Enhanced All Routes Explorer reduction in meshed networks (page 6-7)
- General explorer filters for token ring networks (page 6-9)
- NetBIOS filters (page 6-10)
- SNA filters (page 6-12)
- IPX filters (page 6-13)
- IP filters (page 6-15)
- Other protocol filters (page 6-17)
- Monitor and additional destination filters (page 6-18)
- Advanced filters for token ring (page 6-20)
- Outbound filters (page 6-26)

Enhanced All Routes Explorer Reduction in Meshed Networks

The Centillion 50/100 and Model 5000BH switches include excellent built-in explorer reduction mechanisms that operate automatically. If desired, the CN_BN filter is an additional built-in filter that can further reduce the number of All Routes Explorer (ARE) frames in heavily redundant token ring switch networks. This filter is designed to discard ARE frames that have already traveled through a redundant switch by checking the contents of the frame's Routing Information Field (RIF) for the source routing bridge number of the other switch.

The CN_BN filter is different from other user-defined packet filters, because the only field that has to be supplied by the user is the Name field. In the Name field, you need to enter:

```
CN_BNx
```

where x is the source route bridge number of the *other* redundant switch/bridge.

This bridge number should already be a known and unique value because redundant/parallel source routing bridges require unique bridge numbers. The other fields can be left with the defaults given by the management station.

Note: The CN_BNx filter must be the first filter applied to a token ring port filter list.

Figure 6-5 shows an example of using the CN_BN filter.

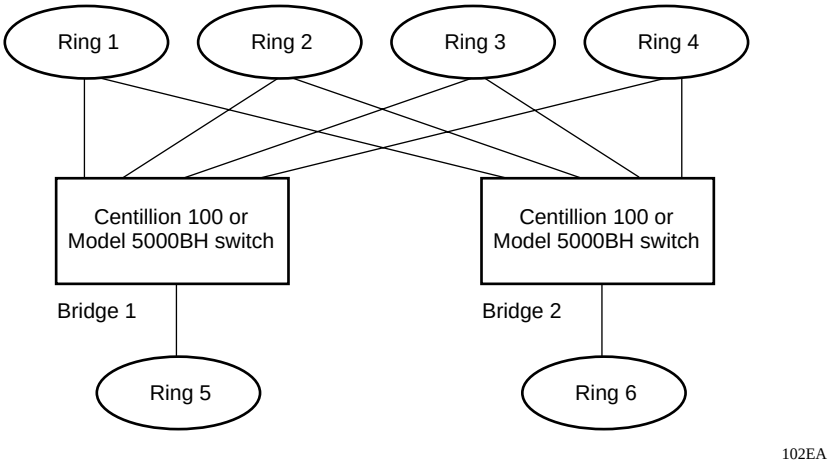

Figure 6-5. Using the CN_BN Filters with Redundant Token Ring Switches

In Figure 6-5, the switch on the left (bridge 1) has the CN_BN2 filter defined (because the number of the other bridge is 2), and it has been applied to the token ring ports that connect to rings 1 through 4. The switch on the right (bridge 2) has the CN_BN1 filter defined, and it has been applied to the token ring ports that connect to rings 1 through 4. Again, note that the CN_BN filter has been defined using the number of the other bridge.

With the CN_BN filter applied in this manner, an ARE frame generated on ring 4 that is subsequently placed on ring 1 by bridge 2 will be discarded by bridge 1 when the ARE frame is received from ring 1. Without the application of the CN_BN filter, bridge 1 would copy the ARE frame out to rings 2 and 3.

General Explorer Filters for Token Ring Networks

This section provides examples of general explorer filters for token ring ports.

Dropping All Routes and Spanning Tree Explorers

Filter FLTR_A in Figure 6-6 drops all routes and spanning tree explorers from a port.

Name	ID	Seq	Type	Offset	Value (hex)	Cond	Match	Fail	Fwd	Mon Dest	Addl Dests
FLTR_A	1	1	MAC	8	80	GE	2	255	NORM		
		2	MAC	14	80	GE	0	0	DROP		

Figure 6-6. Filter to Drop All Routes and Spanning Tree Explorers

FLTR_A sequence 1 is a MAC filter condition with an offset of 8 bytes, which is the beginning of the source address field. This sequence checks if the source routing indicator is turned on or off.

If the first byte of the source address is greater than or equal to 80, the frame continues to sequence 2. If the first byte is less than 80, the frame skips sequence 2 and goes to the next filter group applied on the port. If no other filter groups are applied to the port, the frame is forwarded as normal.

Sequence 2, also a MAC filter, offsets 14 bytes into the frame to check the broadcast indicators in the routing control field. If the first byte is greater than or equal to 80, the frame is dropped and does not continue to the next filter group applied to the port. If the match condition fails, the frame ends the filtering process for that port and is forwarded as normal.

Dropping Spanning Tree Explorers

Filter FLTR_B in Figure 6-7 drops spanning tree explorers from a port.

Name	ID	Seq	Type	Offset	Value (hex)	Cond	Match	Fail	Fwd	Mon Dest	Addl Dests
FLTR_B	1	1	MAC	8	80	GE	2	255	NORM		
		2	MAC	14	C0	GE	0	255	DROP		

Figure 6-7. Filter to Drop Spanning Tree Explorers

FLTR_B sequence 1 with MAC offset of 8 bytes checks for the routing indicator bit. If the routing indicator is turned on, the frame continues to sequence 2. Sequence 2 checks if the frame is a spanning tree explorer.

If the routing indicator is not turned on, the frame continues to the next filter group applied to the port. If there are no other applied filter groups, the frame is forwarded as normal.

Sequence 2 offsets 14 bytes into a frame to check if the frame is a spanning tree explorer. If the filter matches, the frame is dropped. If the frame is not a spanning tree explorer, the frame goes to the next filter group (if any).

NetBIOS Filters

This section provides examples of NetBIOS filters for Ethernet and token ring ports.

Dropping NetBIOS Frames (802.3 Format) from Ethernet Ports

Filter FLE_A in Figure 6-8 drops 802.3 NetBIOS frames from an Ethernet port.

Name	ID	Seq	Type	Offset	Value (hex)	Cond	Match	Fail	Fwd	Mon Dest	Addl Dests
FLE_A	1	1	MAC	14	F0	EQ	0	255	DROP		

Figure 6-8. Filter to Drop NetBIOS Frames (802.3) from an Ethernet Port

FLE_A sequence 1 specifies a MAC offset of 14 bytes to read the DSAP in an 802.3 Ethernet frame. If the DSAP is equal to F0, the frame is dropped and does not continue to the next applied filter on the port. If the DSAP is not equal to F0, the frame is forwarded and continues to the next applied filter.

Dropping NetBIOS Frames from Token Ring Ports

Filter FLTR_D in Figure 6-9 drops NetBIOS frames from a token ring port.

Name	ID	Seq	Type	Offset	Value (hex)	Cond	Match	Fail	Fwd	Mon Dest	Addl Dests
FLTR_D	1	1	LLC	0	F0	EQ	0	255	DROP		

Figure 6-9. Filter to Drop NetBIOS Frames from a Token Ring Port

FLTR_D sequence 1 is an LLC offset of 0 bytes, which points to the DSAP field of a token ring frame. If the DSAP is equal to F0 (which is the SAP for NetBIOS), the frame is dropped. If the match fails, the frame proceeds to the next applied filter group on the port.

Dropping Single Route and All Routes NetBIOS Explorers from Token Ring Ports

Filter FLTR_E in Figure 6-10 drops single and all routes NetBIOS explorers from a token ring port.

Name	ID	Seq	Type	Offset	Value (hex)	Cond	Match	Fail	Fwd	Mon Dest	Addl Dests
FLTR_E	1	1	LLC	0	F0	EQ	2	255	NORM		
		2	MAC	8	80	GE	3	255	NORM		
		3	MAC	14	80	GE	0	255	DROP		

Figure 6-10. Filter to Drop Single and All Routes NetBIOS Explorers

FLTR_E sequence 1 looks at the DSAP to determine if the frame is a NetBIOS frame. If the value is equal to F0, which designates NetBIOS, the frame proceeds to sequence 2. If the value is not equal to 80, the frame goes to the next applied filter group on the interface.

Sequence 2 is a MAC offset of 8 bytes that checks for the routing indicator bit. If the routing indicator bit is turned on, the frame continues to sequence 3. If not, the frame goes to the next applied filter group on the port.

Sequence 3 checks for explorers; if there is a match, the frame is dropped and does not continue to any other filter groups. If the frame is not an explorer, it proceeds to the next applied filter group. If there are no other applied filter groups, the frame is forwarded as normal.

SNA Filters

This section provides examples of SNA filters for token ring and Ethernet ports.

Forwarding SNA Frames and Dropping All Other Frames from Token Ring Ports

Filter FLTR_F in Figure 6-11 forwards SNA frames and drops all others.

Name	ID	Seq	Type	Offset	Value (hex)	Cond	Match	Fail	Fwd	Mon Dest	Addl Dests
FLTR_F	1	1	LLC	0	00	NE	2	0	DROP		
		2	LLC	0	04	EQ	0	0	NORM		

Figure 6-11. Filter to Forward Only SNA Frames on Token Ring Ports

Sequence 1 matches on all frames except on a null DSAP of 00, which is commonly used by SNA Test Poll frames to discover the host. If a frame equals DSAP 00, sequence 1 fails and the frame is forwarded as normal. Otherwise, the new frame handling state is set to DROP from the initial state of NORM.

If sequence 2 matches SNA DSAP of 04, the frame is forwarded normally.
If sequence 2 fails (the frame does not have DSAP of 04), the filter stops and the previous frame handling state is used, which in this case is DROP the frame.

Forwarding SNA Frames (802.3) and Dropping All Other Frames from Ethernet Ports

Filter FLE_B in Figure 6-12 forwards 802.3 SNA frames and drops everything else from Ethernet ports.

Name	ID	Seq	Type	Offset	Value (hex)	Cond	Match	Fail	Fwd	Mon Dest	Addl Dests
FLE_B	1	1	MAC	14	00	NE	2	0	DROP		
		2	MAC	14	04	EQ	0	0	NORM		

Figure 6-12. Filter to Forward Only SNA Frames (802.3) on Ethernet Ports

In sequence 1, the DSAP field in an 802.3 Ethernet frame is checked for null DSAP. If the DSAP is not equal to 00, the frame goes to sequence 2 and the frame Forwarding state is set to DROP. If it is equal to 00 (the condition in sequence 1 fails), the frame is forwarded.

Sequence 2 checks for the SNA DSAP of 04. If this condition matches, the frame is forwarded and the Forwarding state is reset to NORM. If the DSAP is not equal to 04, the frame is dropped.

Forwarding SNA Frames (Ethernet V2.0) and Dropping All Other Frames from Ethernet Ports

Filter FLE_C in Figure 6-13 forwards version 2.0 SNA frames and drops everything else from Ethernet ports.

Name	ID	Seq	Type	Offset	Value (hex)	Cond	Match	Fail	Fwd	Mon Dest	Addl Dests
FLE_C	1	1	MAC	12	80.D5	NE	0	0	DROP		

Figure 6-13. Filter to Forward Only SNA Frames (2.0) on Ethernet Ports

Sequence 1 checks if the protocol type field is equal to 80d5, which designates IBM SNA services over Ethernet. If the type field is not equal to 80d5, the frame is dropped and does not continue to the next applied filter. The Forwarding state is set to DROP. If the type field is equal to 80d5, the frame is forwarded and proceeds to the next applied filter group on the port.

IPX Filters

This section provides examples of IPX filters for token ring and Ethernet ports.

Dropping Token Ring IPX Traffic and Forwarding All Others

Filter FLTR_G in Figure 6-14 drops IPX traffic and forwards all other frames on token ring ports.

Name	ID	Seq	Type	Offset	Value (hex)	Cond	Match	Fail	Fwd	Mon Dest	Addl Dests
FLTR_G	1	1	LLC	0	E0	EQ	0	255	DROP		

Figure 6-14. Filter to Drop IPX Traffic on Token Ring Ports

FLTR_G sequence 1 indicates that the filter offset type is LLC with no byte offset, which points to the DSAP field of a token ring packet. Any frame with a DSAP equal to E0 is dropped. The frame does not continue to the next applied filter group on the port. The Forwarding state is set to DROP.

If the DSAP does not equal E0, the frame is forwarded and proceeds to the next applied filter group.

Dropping IPX Frames and Forwarding All Others on Ethernet Ports

Filters FLE_D and FLE_E in Figure 6-15 drop all version 2.0 IPX frames and forward all others on Ethernet ports.

Name	ID	Seq	Type	Offset	Value (hex)	Cond	Match	Fail	Fwd	Mon Dest	Addl Dests
FLE_D	1	1	MAC	12	81.37	EQ	0	255	DROP		
FLE_E	2	1	MAC	12	81.38	EQ	0	255	DROP		

Figure 6-15. Filter to Drop IPX Frames (V2.0) on Ethernet Ports

Assume that both FLE_D and FLE_E are applied to the same port. Two filters are needed to check for IPX protocol types because Novell uses both 8137 and 8138.

FLE_D sequence 1 checks for the protocol type field for 8137. If there is a match, the frame is dropped and the Forwarding state is set to DROP. The DROP Forwarding state on FLE_D ensures that the next applied filter, FLE_E, is not carried out.

If the type field is not 8137, the next applied filter group is checked and the frame Forwarding state remains at NORM. If the type field is 8138, the frame is dropped. If it is not 8138, the frame is forwarded.

IP Filters

This section provides examples of IP filters on token ring and Ethernet ports.

Dropping IP Frames (Not Including ARP and RARP) from Token Ring Ports

Filter FLTR_H in Figure 6-16 drops all IP frames (except ARP and RARP) from token ring ports.

Name	ID	Seq	Type	Offset	Value (hex)	Cond	Match	Fail	Fwd	Mon Dest	Addl Dests
FLTR_H	1	1	LLC	6	08.00	EQ	0	255	DROP		

Figure 6-16. Filter to Drop All IP Frames Except ARP and RARP from Token Ring Ports

FLTR_H sequence 1 is an LLC filter with offset of 6 bytes, which points to the type field in the SNAP portion of an IP token ring frame. If the value matches, the filter drops the frame. If the value does not match, the frame is forwarded to the next filter group on the port.

Dropping Version 2.0 IP Frames (Not Including ARP and RARP) from Ethernet Ports

Filter FLE_F in Figure 6-17 drops all version 2.0 IP frames (except ARP and RARP frames) from Ethernet ports.

Name	ID	Seq	Type	Offset	Value (hex)	Cond	Match	Fail	Fwd	Mon Dest	Addl Dests
FLE_F	1	1	MAC	12	08.00	EQ	0	255	DROP		

Figure 6-17. Filter to Drop All IP Frames (V2.0) Except ARP and RARP from Ethernet Ports

FLE_F sequence 1 offsets 12 bytes to the protocol type field in an Ethernet V2.0 frame. If the type field is equal to 0800, the frame is dropped and does not continue to the next applied filter group. If the type field is not equal to 0800, the frame is forwarded.

Dropping ARP and RARP Frames from Token Ring Ports

Filter FLTR_I in Figure 6-18 drops ARP and RARP frames from token ring ports.

Name	ID	Seq	Type	Offset	Value (hex)	Cond	Match	Fail	Fwd	Mon Dest	Addl Dests
FLTR_I	1	1	LLC	6	08.06	EQ	0	2	DROP		
		2	LLC	6	08.35	EQ	0	255	DROP		

Figure 6-18. Filter to Drop ARP and RARP Frames from Token Ring Ports

FLTR_I uses an LLC offset of 6 bytes to match the type field in the SNAP header. If the type field is equal to 0806, which designates ARP, the frame is dropped and the frame handling state is set to DROP.

If the frame is not an ARP, the frame goes to sequence 2 to check if the frame is RARP. If the frame is RARP, the frame is dropped. If not, the frame proceeds to the next filter group. If no other filter groups are applied on the port, the frame is forwarded as normal.

Dropping ARP and RARP Frames from Ethernet Ports

Filter FLE_G in Figure 6-19 drops ARP and RARP frames from Ethernet ports.

Name	ID	Seq	Type	Offset	Value (hex)	Cond	Match	Fail	Fwd	Mon Dest	Addl Dests
FLE_G	1	1	MAC	12	08.06	EQ	0	2	DROP		
		2	MAC	12	08.35	EQ	0	255	DROP		

Figure 6-19. Filter to Drop ARP and RARP Frames from Ethernet Ports

FLE_G sequence 1 checks if the protocol type field in an Ethernet V2.0 frame is equal to 0806, which designates ARP frames. If there is a match, the frame is dropped and the next applied filter on the port is not carried out.

If the match fails, sequence 2 is carried out, which checks for RARP frames. If there is a match, the RARP frame is dropped and the filter stops. If the frame is neither ARP nor RARP, it is forwarded.

Other Protocol Filters

This section provides examples of the following filter applications:

- Dropping Banyan frames from token ring ports
- Dropping DEC LAT frames from Ethernet ports

Dropping Banyan Frames from Token Ring Ports

Filter FLTR_J in Figure 6-20 drops Banyan frames from token ring ports.

Name	ID	Seq	Type	Offset	Value (hex)	Cond	Match	Fail	Fwd	Mon Dest	Addl Dests
FLTR_J	1	1	LLC	6	0B.AD	EQ	0	255	DROP		

Figure 6-20. Filter to Drop Banyan Frames from Token Ring Ports

FLTR_J drops all LLC frames with SNAP type of 0BAD, which is used for Banyan VINES frames.

Dropping DEC LAT Frames from Ethernet Ports

Filter FLE_H in Figure 6-21 drops DEC LAT frames from Ethernet ports.

Name	ID	Seq	Type	Offset	Value (hex)	Cond	Match	Fail	Fwd	Mon Dest	Addl Dests
FLE_H	1	1	MAC	12	60.04	EQ	0	255	DROP		

Figure 6-21. Filter to Drop DEC LAT Frames from Ethernet Ports

Sequence 1 offsets 12 bytes into Ethernet V2.0 frames to check the protocol type field. If the type field equals 6004, which is commonly used by LAT, the frame is dropped and does not continue to the next applied filter. If the type field is not equal to 6004, the frame is forwarded and continues to the next applied filter group on the port.

Monitor and Additional Destination Filters

A filter can specify any LAN port as a monitor port. The monitor port (Mon Dest) should be connected to a statistics gathering tool such as an RMON probe or a network analyzer. The additional destination port or ports (Addl Dest) are used to manually redirect frames to specified ports.

The Centillion architecture does not support translational bridging between Ethernet and token ring; therefore, token ring traffic cannot be monitored on an Ethernet port or vice versa.

Two forwarding actions can be used with the monitor and additional destination assignments. NORM forwards a copy of the frame to the specific monitor destination or additional destination as well as to the intended destination. ALT forwards a copy of the frame to the specified monitor or additional destination only, without forwarding to the intended destination.

Monitoring All Frames on Token Ring Ports

Filter FLTR_K in Figure 6-22 sends all frames on token ring ports to a monitor destination.

Name	ID	Seq	Type	Offset	Value (hex)	Cond	Match	Fail	Fwd	Mon Dest	Addl Dests
FLTR_K	1	1	MAC	2	00	GE	0	255	NORM	3:1	

Figure 6-22. Filter to Send Frames on Token Ring Ports to a Monitor Destination

All frames with the first byte of destination address greater than or equal to 00 are monitored on switch module 3 port 1(3:1). The NORM action specifies that all matched frames are mirrored to the monitor port and also forwarded to the normal destination. If the desired action is to mirror the frames and not forward to normal destination, use ALT as the forwarding action instead of NORM.

Monitoring Frames on Ethernet Ports

Filter FLE_I in Figure 6-23 sends all frames from Ethernet ports to a monitor destination.

Name	ID	Seq	Type	Offset	Value (hex)	Cond	Match	Fail	Fwd	Mon Dest	Addl Dests
FLE_I	1	1	MAC	0	00	GE	0	255	NORM	3:16	

Figure 6-23. Filter to Send Frames from Ethernet Ports to a Monitor Destination

All frames with the first byte of destination address greater than or equal to 00 are mirrored on card 3, port 16; they are also forwarded to their normal destinations.

Forwarding Frames to Specific Token Ring Ports

Filter FLTR_L in Figure 6-24 forwards frames to specific ports on token ring modules.

Name	ID	Seq	Type	Offset	Value (hex)	Cond	Match	Fail	Fwd	Mon Dest	Addl Dests
FLTR_L	1	1	MAC	2	00	GE	0	255	ALT		2:3,4:4

Figure 6-24. Filter to Forward Frames to Specific Token Ring Ports

This filter specifies that all frames with the first byte of destination address greater than or equal to 00 are forwarded to switch module 2, port 3 (2:3) and to switch module 4, port 4 (4:4). The forwarding action ALT redirects all frames to the additional destinations without forwarding to the normal destinations.

Advanced Filters for Token Ring

Figure 6-25 shows an example of directing routed protocols to router ports and switching layer 2 locally.

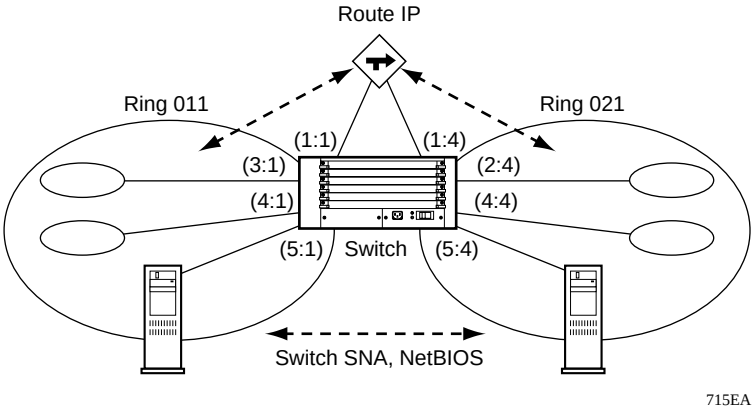

Figure 6-25. Directing Routed Protocols to Router Ports and Switching Layer 2 Locally

In the examples shown in Figure 6-25, virtual ring 011 is configured for ports 1:1, 3:1, 4:1, and 5:1 with the router port on 1:1; virtual ring 021 is configured for ports 1:4, 2:4, 4:4, and 5:4 with the router port on 1:4. The switching mode on the switch is configured for source route bridging.

The objective is to have all IP traffic forwarded between rings 011 and 021 by the router only, while all other traffic (for example, SNA, NetBIOS, IPX) is bridged between rings 011 and 021 by the switch.

The filters shown in Figure 6-26 provide the identification of the IP traffic to be controlled, while all other traffic is treated normally by the switch.

Name	ID	Seq	Type	Offset	Value (hex)	Cond	Match	Fail	Fwd	Mon Dest	Addl Dests
IP1	1	1	MAC	8	80	GE	2	0	NORM		
		2	MAC	14	80	GE	3	0	NORM		
		3	LLC	0	AA.AA.03.00.00.00.08.00	EQ	0	4	ALT		1:1,3:1,4:1,5:1
		4	LLC	0	AA.AA.03.00.00.00.08.06	EQ	0	0	ALT		1:1,3:1,4:1,5:1
IP2	2	1	MAC	8	80	GE	2	0	NORM		
		2	MAC	14	80	GE	3	0	NORM		
		3	LLC	0	AA.AA.03.00.00.00.08.00	EQ	0	4	ALT		1:4,2:4,4:4,5:4
		4	LLC	0	AA.AA.03.00.00.00.08.06	EQ	0	0	ALT		1:4,2:4,4:4,5:4

Figure 6-26. Filters to Send Frames on Ethernet Ports to a Monitor Destination

The two filters each identify source route explorer IP frames and restrict them to the ports in the local virtual ring. The IP1 filter is applied to ports 1:1, 3:1, 4:1, and 5:1; the IP2 filter is applied to ports 1:4, 2:4, 4:4, and 5:4.

To ensure that IP frames are forwarded to rings within a virtual ring only, and not switched by the Centillion 100 switch, the flow of source route explorer IP frames must be controlled so that they are forwarded from a ring within a virtual ring to other rings in the same virtual ring only.

It is acceptable for the router also to operate as an SRB, as long as it can intercept the IP source route explorers without forwarding them. This feature allows the router to act as an SRB for wide area bridging. Filters should be installed on the router to prevent it from forwarding explorers between virtual rings 011 and 021.

The sequences in IP1 and IP2 filters work as follows:

- Sequence 1: Examines the routing indicator of each frame to determine if it is source routed; if not, treat the frame normally.
- Sequence 2: Examines the route control to determine if it is an explorer; if not, treat the frame normally.
- Sequence 3: Examines the type field in the IP SNAP header; if it is IP, forward the frame only to the ports of the virtual ring. Note that the Centillion 100 switch never forwards a frame to the port from which it was received, so that the inclusion of the source port in the “Addl Dest” is not a conflict.
- Sequence 4: Examines the frame content for the IP ARP SNAP header. If it is IP ARP, forward the frame only to the ports of the virtual ring.

Restricting IPX Get Nearest Server Broadcasts to the Local Ring

In the case where all servers accessed from the work group ring are located on the local ring, a filter to drop all Get Nearest Server (GNS) requests can be applied on the Centillion 100 switch to reduce the broadcasts sent to the network (Figure 6-27).

Name	ID	Seq	Type	Offset	Value (hex)	Cond	Match	Fail	Fwd	Mon Dest	Addl Dests
IPX1	1	1	MAC	2	FF.FF.FF.FF.FF.FF	EQ	2	0	NORM		
		2	LLC	0	E0.E0.03	EQ	255	0	NORM		
GNS1	2	1	LLC	13	FF.FF.FF.FF.FF.FF.04.52	EQ	2	255	NORM		
		2	LLC	33	00.03	EQ	0	3	DROP		
		3	LLC	33	00.01.00.04	EQ	0	255	DROP		

Figure 6-27. Filters to Restrict IPX GNS Broadcasts to the Local Ring

Filter IPX1 screens for IPX broadcast frames only and must be applied on an interface first before GNS1 is applied. If the frame is an IPX broadcast, filter GNS1 checks if the broadcast is GNS.

Note: Sequence 3 is optional. Sequence 3 of GNS1 is required if VLMs are used in the client machines and you want to restrict them to the local area even if GNS fails. This filter without sequence 3 for VLM-equipped clients ensures that a local server is used if available, and then it finds any server in the network if the local GNS fails, because the client then needs to issue a Get General File Server request.

Restricting IPX Get Nearest Server Broadcasts to Workgroup Rings

A filter can be constructed to forward all GNS requests generated on the ring to only the nearby rings containing the requested servers (Figure 6-28).

Name	ID	Seq	Type	Offset	Value (hex)	Cond	Match	Fail	Fwd	Mon Dest	Addl Dests
IPX1	1	1	MAC	2	FF.FF.FF.FF.FF.FF	EQ	2	0	NORM		
		2	LLC	0	E0.E0.03	EQ	255	0	NORM		
GNS1	2	1	LLC	13	FF.FF.FF.FF.FF.FF.04.52	EQ	2	255	NORM		
		2	LLC	33	00.03	EQ	0	255	ALT		1:3,2:4

Figure 6-28. Filter to Restrict IPX GNS Broadcasts to Workgroup Rings

The sequences work as follows:

- IPX1Sequence 1: Checks if the frame is an all Fs MAC broadcast. If it is, go to sequence 2. If not, then forward the frame and do not go on to the next applied filter on the port.
- IPX1 Sequence 2: Checks if the frame is IPX GNS. If it is, go to the next filter; if it is not an IPX GNS, then stop filtering and do not go to the next applied filter on the port.
- GNS2 Sequence 1: Checks if the IPX address is an all Fs broadcast. If it is, go to 2; if not, forward the frame and go to the next applied filter on the port.
- GNS2 Sequence 2: Checks if the frame is a GNS. If it is, forward frames to the specified destinations under Addl Dest.

Restricting RIP and SAP Broadcasts to the Local Ring

The filter shown in Figure 6-29 identifies RIP and SAP updates and drops them before they are propagated into the rest of the network. The filter is applied to the Centillion 50/100 or the Model 5000BH switch port connecting the local ring.

Name	ID	Seq	Type	Offset	Value (hex)	Cond	Match	Fail	Fwd	Mon Dest	Addl Dests
IPX1	1	1	MAC	2	FF.FF.FF.FF.FF.FF	EQ	2	0	NORM		
		2	LLC	0	E0.E0.03	EQ	255	0	NORM		
GNS2	2	1	LLC	13	FF.FF.FF.FF.FF.FF.04.52	EQ	3	2	NORM		
		2	LLC	13	FF.FF.FF.FF.FF.FF.04.53	EQ	3	255	NORM		
		3	LLC	33	00.02	EQ	0	255	DROP		

Figure 6-29. Filter to Restrict RIP and SAP Broadcasts to the Local Ring

Restricting RIP and SAP Broadcasts to within the Rings for a Workgroup

The filter shown in Figure 6-30 identifies RIP and SAP broadcasts and forwards them only to the other rings in the work group.

Name	ID	Seq	Type	Offset	Value (hex)	Cond	Match	Fail	Fwd	Mon Dest	Addl Dests
IPX1	1	1	MAC	2	FF.FF.FF.FF.FF.FF	EQ	2	0	NORM		
		2	LLC	0	E0.E0.03	EQ	255	0	NORM		
GNS2	2	1	LLC	13	FF.FF.FF.FF.FF.FF.04.52	EQ	3	2	NORM		
		2	LLC	13	FF.FF.FF.FF.FF.FF.04.53	EQ	3	255	NORM		
		3	LLC	33	00.02	EQ	0	255	ALT		3:2,4:1

Figure 6-30. Filter to Restrict RIP and SAP Broadcasts to Within the Rings for a Workgroup

Reducing HP Laser Jet Printer Get General File Server Broadcasts

When an HP Laser Jet printer loses contact with the file server, it sends three Get General File Server broadcasts, one second apart, every 10 seconds. To confine these broadcasts to the local ring, the filter shown in Figure 6-31 drops these broadcasts. This filter should be applied on switch ports that should not forward the Get General File Server broadcasts to other rings.

Name	ID	Seq	Type	Offset	Value (hex)	Cond	Match	Fail	Fwd	Mon Dest	Addl Dests
IPX1	1	1	MAC	2	FF.FF.FF.FF.FF.FF	EQ	2	0	NORM		
		2	LLC	0	E0.E0.03	EQ	255	0	NORM		
GNS2	2	1	LLC	13	FF.FF.FF.FF.FF.FF.04.52	EQ	2	255	NORM		
		2	LLC	8	90.00.90	EQ	3	255	NORM		
		3	LLC	33	00.01.00.04	EQ	0	255	DROP		

Figure 6-31. Filter to Reduce HP Laser Jet Printer Get General File Server Broadcasts

Restricting HP Laser Jet Printer Get General File Server Broadcasts to a Workgroup

The filter shown in Figure 6-32 screens for HP Laser Jet Get General File Server broadcasts and forwards them to the specified ports.

Name	ID	Seq	Type	Offset	Value (hex)	Cond	Match	Fail	Fwd	Mon Dest	Addl Dests
IPX1	1	1	MAC	2	FF.FF.FF.FF.FF.FF	EQ	2	0	NORM		
		2	LLC	0	E0.E0.03	EQ	255	0	NORM		
GNS2	2	1	LLC	13	FF.FF.FF.FF.FF.FF.04.52	EQ	2	255	NORM		
		2	LLC	8	90.00.90	EQ	3	255	NORM		
		3	LLC	33	00.01.00.04	EQ	0	255	ALT		4:1,3:5

Figure 6-32. Filter to Restrict HP Laser Jet Get General File Server Broadcasts to a Workgroup

Outbound Filters

Although the filter examples in "Packet Filter Examples" can apply to either inbound or outbound filters, some outbound filters are more helpful in particular situations.

Routers use outbound filters for prioritization of traffic that is only supported on WAN interfaces. Switches my use outbound filters to protect a link to an Ethernet attached router. A network administrator may not be worried about incoming broadcasts to a switch but may not want to send the broadcast packets to the router or another non-Centillion or non-Bay Networks switch. In such a case, the administrator would create a standard packet filter for a broadcast packet and apply it to the port as an outbound filter.

Outbound Filter Examples

This section provides outbound filter examples that may be useful in your network.

Copying All Outgoing IP Packets from an Ethernet Port to a Monitor Destination

The filter in Figure 6-33 copies all outgoing IP packets from a targeted Ethernet port to a monitor destination on port 1 of module 1.

Name	ID	Seq	Type	Offset	Value (hex)	Cond	Match	Fail	Fwd	Mon Dest	Addl Dests
IP_F	1	1	MAC	12	08.00	EQ	0	255	NORM		

Figure 6-33. Filter to Copy Outgoing IP Packets to a Monitor Destination

Dropping All Outgoing UDP Packets from Ethernet Ports

The filter in Figure 6-34 drops all outgoing UDP packets form Ethernet ports.

Name	ID	Seq	Type	Offset	Value (hex)	Cond	Match	Fail	Fwd	Mon Dest	Addl Dests
Udp_D	1	1	MAC	23	11	EQ	0	255	DROP		

Figure 6-34. Filter to Drop All Outgoing UDP Packets from Ethernet Ports

Dropping All Outgoing PING (ICMP) Packets from Ethernet Ports

The filter in Figure 6-35 drops all outgoing PING (ICMP) packets from Ethernet ports.

Name	ID	Seq	Type	Offset	Value (hex)	Cond	Match	Fail	Fwd	Mon Dest	Addl Dests
Ping_d	1	1	MAC	23	01	EQ	0	255	DROP		

Figure 6-35. Filter to Drop All Outgoing PING (ICMP) Packets from Ethernet Ports

Common SAPs

Table 6-1 lists common service access points (SAPs) that may be useful when designing packet filter groups.

Table 6-1. Common SAPs

SAP	Assignment
00	Null LSAP
42	IEEE 802.1 Bridge Spanning Tree Protocol
AA	Sub-Network Access Protocol (SNAP)
04	IBM SNA Path Control (Individual)
05	IBM SNA Path Control (Group)
18	Texas Instruments
80	Xerox Network Systems (XNS)
BC	Banyan VINES
E0	Novell NetWare
F0	IBM NetBIOS
F4	IBM LAN Management (Individual)
F5	IBM LAN Management (Group)
F8	IBM Remote Program Load (RPL)
FA	Ungermann Bass

Ethernet Protocol Types

Table 6-2 lists some common Ethernet protocol types that may be useful when designing packet filter groups.

Table 6-2. Ethernet Protocol Types

Value	Description
0600	Xerox XNS IDP
0800	DOD IP
0805	X.25 Level 3
0806	IP ARP
0BAD	Banyan Systems
6001	DEC MOP Dump/Load Assistance
6002	DEC MOP Remote Console
6003	DEC DECnet Phase IV
6004	DEC LAT
6005	DEC DECnet Diagnostics
6010–6014	3Com Corporation
7000–7002	Ungermann-Bass Download
7030	Proteon
7034	Cabletron
8035	IP Reverse ARP
8046–8047	AT&T
8088–808A	Xyplex
80C0–80C3	Digital Communication Associates
80D5	IBM SNA Services over Ethernet
80FF–8103	Bay Networks Wellfleet
8137–8138	Novell

Chapter 7
NetBIOS Protocol

The NetBIOS protocol uses 16-character station names and a broadcast explorer route discovery mechanism to establish sessions with its peers or servers. In a complex bridged network, the large number of discovery frames uses network bandwidth and can cause bridge congestion by duplicating broadcast explorer frames. The switch reduces broadcast explorer frames by learning NetBIOS names and then directly switching to the destination port.

Certain NetBIOS frames are affected by name filters, name proxy, and broadcast discard. These frames are transmitted by NetBIOS stations as UI frames to the NetBIOS functional address. When any one of the three NetBIOS features is set, the internal filter is set to receive these frames and hand them to the MCP for NetBIOS processing. NetBIOS user data frames are transmitted to a Unicast address in SRF (Specifically Routed Frame) format and are not affected by these NetBIOS features. Existing NetBIOS sessions are not disturbed.

Note: The order of NetBIOS frame processing is to apply the DATAGRAM_BROADCAST discard first, then NetBIOS Name Filters, and then NetBIOS Name Proxy.

Table 7-1 shows the NetBIOS features and corresponding frame types that are checked by each feature.

Table 7-1. NetBIOS Features and Frame Types

Feature	NetBIOS Frame Type	Description
Name filters	ADD_GROUP_NAME_QUERY	This frame is used to verify that the group name to be added is not already used as a unique name within the network. The sender's name is checked against the supplied filter name.
	ADD_NAME_QUERY	This frame is used to verify that the name to be added is not already used as a unique name within the network. The sender's name is checked against the supplied filter name.
	NAME_QUERY	This frame is used to request that a remote node establish a session or to locate a name on the network. The recipient's name is checked against the supplied filter name.
	DATAGRAM	This frame is used to transmit user data to one name in the network. The destination name is the name of the recipient. The recipient's name is checked against the supplied filter name.
Name proxy	NAME_QUERY	This frame is used to request that a remote node establish a session or to locate a name in the network.
	STATUS_QUERY	This frame is used to request remote adapter status.
	DATAGRAM	This frame is used to transmit user data to one name in the network.

Name Filters and Examples

The NetBIOS Name Filters feature allows the switch to forward or discard certain NetBIOS frame types based on the name of a station. NetBIOS name filtering is on inbound ports. Multiple name filters can be assigned to each port or ring number.

Four NetBIOS frames are affected by these NetBIOS name filters. The configured filter name is applied differently for each frame type. The filter name is compared against the *Source NetBIOS name* for the ADD_GROUP_NAME_QUERY and ADD_NAME_QUERY frames and then compared against *the Destination NetBIOS name* for the NAME_QUERY and DATAGRAM frames.

Note: The sixteenth character of a NetBIOS name has special operational meaning and should not be used when NetBIOS name filters are configured.

NetBIOS names can contain the question mark (?) and/or asterisk (*) wildcard characters. The? character can be anywhere in the name, and the character in that position is ignored for comparison. For example, a name SC0?_01 will match SC01_01 through SC09_01. The * character must be the last character in the name. For example, SC01_* will match any name that begins with SC01_. The action specified (Forward or Discard) is taken if a match is found; otherwise, the opposite action is taken.

For example, if a name SC01_01 is configured with Discard action:

- ADD_GROUP_NAME_QUERY and ADD_NAME_QUERY frames with source NetBIOS names that are SC01_01 are dropped.
- NAME_QUERY and DATAGRAM frames with destination NetBIOS names that are SC01_01 are dropped.
- Other types of NetBIOS frames and frames that do not match SC01_01 are switched normally according to the configured bridging mode.

For example, if a filter is on module 4, port 2, with the name Station_A and the action Discard, then any ADD_GROUP_NAME_QUERY or ADD_NAME_QUERY frames from Station_A will be discarded.

Name or status queries are used when a station is trying to find a particular recipient. The station sends out either a NAME_QUERY or STATUS_QUERY frame. If a NetBIOS name filter has been assigned to the port to which the station is attached, these frames are checked for the recipient's name and forwarded or discarded based on the assigned action.

For example, if a filter is on module 4, port 2, with the name Station_B and the action Discard, any NAME_QUERY or STATUS_QUERY frames for station B are discarded. Station A is not able to establish subsequent NetBIOS sessions with station B; previously established sessions are not affected.

The network configuration for the examples in this section is shown in Figure 7-1.

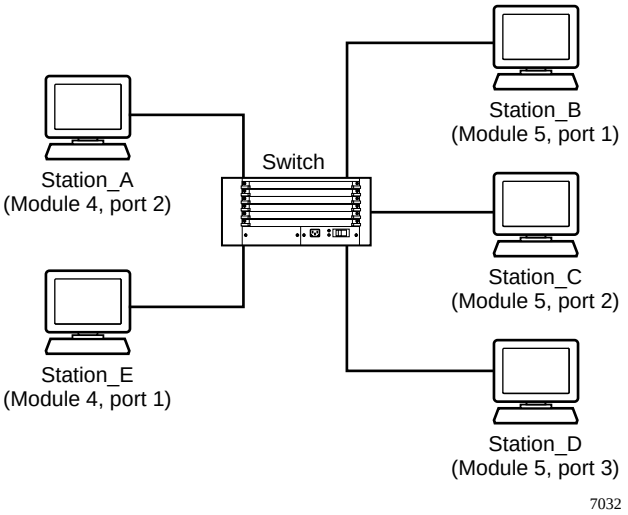

Figure 7-1. NetBIOS Filter Assignments

If you enter a NetBIOS filter on module 4, port 2, specifying * as the name and Discard as the action, you would stop NetBIOS traffic coming from station A.

If you put a NetBIOS filter on module 5, ports 1, 2, and 3, specifying the name Station_A and the action Discard, you would stop NetBIOS query traffic destined for station A.

In addition, you could keep station B's NetBIOS traffic from going to station C with the name Station_B and action Discard on module 5, port 2.

On station E, you could only allow NetBIOS traffic to station C by specifying the name Station_C and action “Forward” on module 4, port 1.

Note: Ring numbers are used instead of the module number, port number combination when source route bridging or source route transparent bridging are used as the switching mode.

This sample data is summarized in Table 7-2.

Table 7-2. NetBIOS Sample Data

Ring Number or Module:Port	Name	Action
4:2	*	Discard
5:1	Station_A	Discard
5:2	Station_A	Discard
5:3	Station_A	Discard
5:2	Station_B	Discard
4:1	Station_C	Forward

When a NetBIOS station is reset or powered up, that station sends multiple ADD_GROUP_NAME_QUERY and ADD_NAME_QUERY frames to establish the uniqueness of the names it intends to use. These two types of frames are transmitted as a UI frame to the NetBIOS functional address using single-route broadcast. If there is a conflict, however, the frame NAME_IN_CONFLICT (0x02) is returned with the sender’s unicast MAC address using a specifically routed RIF. If NetBIOS names are learned from these frames, all NetBIOS frames must be checked for name conflicts, degrading switch performance. Therefore, the switch does not learn NetBIOS names from ADD_GROUPU_NAME_QUERRY and ADD_NAME_QUERY.

NetBIOS Name Proxy is composed of two parts: one learns NetBIOS station names; the second looks up destination names in the NetBIOS Name Table and switches to the appropriate ports. If the destination name is not found in the table, queries are bridged according to the configured bridging mode.

Note: NetBIOS names have to be learned before queries to those names are proxied. Therefore, to test NetBIOS Name Proxy, a NAME_QUERY, DATAGRAM, or STATUS_QUERY frame must be transmitted from the NetBIOS station before any queries to that station are forwarded.

Appendix A
Switch Module and Software Matrix

This appendix contains tables that describe the Centillion 50/100 and Model 5000BH switch host and MPC modules and the minimum software requirement for each. It also includes a switch image compatibility matrix for MCP modules and Centillion software releases 2.1 and greater.

Switch Modules and Minimum Software Requirements

Table A-1 lists all the switch modules for the Centillion 50/100 switch. Table A-2 lists all the switch modules for the Model 5000BH switch.

The modules are arranged by networking technology: Ethernet, token ring, and ATM. The MCP modules are listed first, then the host modules, by each networking technology.

Table A-1. Centillion 50/100 Switch Module and Minimum Required Software Matrix

Part Number	Module Description	Special Notes	Minimum Software Requirement
Centillion 50/100 EtherSpeed Switch Modules			
N/A	8-port EtherSpeed 10BASE-T MCP switch module	Discontinued	2.0.5 only
AS1904001	16-port EtherSpeed 10BASE-T switch module (1.25 MB buffer memory, 1 KB CAM)		2.0
AS1904002	16-port 10BASE-T EtherSpeed switch module (1.25 MB buffer memory, 8 KB CAM)		2.0
AS1904005	8-port 10BASE-FL EtherSpeed switch module		2.0

Table A-1. Centillion 50/100 Switch Module and Minimum Required Software Matrix (continued)

Part Number	Module Description	Special Notes	Minimum Software Requirement
AS1904008	24-port 10BASE-T EtherSpeed switch module (two 50-pin Telco connectors)		3.0.1
AS2004002	10-port 10/100 Mb/s EtherSpeed switch module (eight 10BASE-T ports, two 100BASE-TX ports)		2.0
AS2004004	10-port 10/100 Mb/s EtherSpeed switch module (eight 10BASE-T ports, two 100BASE-FX ports)		2.0
AS3104001	4-port 10/100 Mb/s EtherSpeed switch module		2.1.1
AS3104003	16-port 100 Mb/s EtherSpeed segment switch module (four 4-port 100BASE-TX segments)		2.2.1
Centillion 50/100 TokenSpeed Switch Modules			
N/A	4-port UTP/STP TokenSpeed MCP switch module	Discontinued	1.0
AS1704002	8-port UTP/STP TokenSpeed MCP switch module		2.1.1
AS1104002	4-port MMF TokenSpeed switch module	Discontinued with 4.0	1.0
AS1704001	8-port UTP/STP TokenSpeed switch module		2.1.1
AS1704004	16-port TokenSpeed HD switch module		3.2.1
N/A	4-port UTP/STP TokenSpeed switch module	Discontinued with 4.0	1.0
Centillion 50/100 ATMSpeed Switch Modules and Media Dependent Adapters			
AS1304010	Centillion 50/100 ATMSpeed 622-MM MCP switch module		4.0
AS1304011	Centillion 50/100 ATMSpeed 622-SM MCP switch module		4.0
AS1304008	4-port ATMSpeed MDA MCP switch module (accepts two 2-port CL MDAs)		2.2.1
AS1304009	4-port ATMSpeed MDA MCP switch module (accepts two 2-port CL MDAs)		2.2.1
N/A	4-port wiring closet MMF ATMSpeed/155 MCP switch module		2.1.0
N/A	4-port network center MMF ATMSpeed/155 MCP switch module		2.1.1

Table A-1. Centillion 50/100 Switch Module and Minimum Required Software Matrix (continued)

Part Number	Module Description	Special Notes	Minimum Software Requirement
AS1304012	Centillion 50/100 ATMSpeed 622-MM switch module		4.0.1
N/A	2-port MMF/SMF ATMSpeed/155 ATMSpeed switch module		1.0
CL1333001	2-port SONET/SDH multimode fiber ATMSpeed MDA		2.2.1
CL1333002	2-port SONET/SDH UTP ATMSpeed MDA		2.2.1
CL1333003	2-port SONET/SDH single-mode fiber ATMSpeed MDA		2.2.1
CL1333004	2-port DS-3 ATM MDA		3.2.1
CL1333005	2-port E-3 ATM MDA		3.2.1

Table A-2. Model 5000BH Switch Modules

Part Number	Description	Special Notes	Minimum Software Requirement
Model 5000BH EtherSpeed MCP and Switch Modules			
N/A	Model 5328M 8-port 10BASE-T EtherSpeed MCP switch module	Discontinued	2.0.5 only
N/A	Model 5328 16-port 10BASE-T EtherSpeed switch module		2.0
CL1904003	Model 5324 12-port 10BASE-FL EtherSpeed switch module		2.1.1
CL1904004	Model 5328HD 24-port 10BASE-T EtherSpeed switch module		2.2.1
CL3104004	Model 5425 4-port 10/100 Mb/s EtherSpeed switch module		2.1
CL3104006	Model 5455 16-port 100 Mb/s EtherSpeed segment switch (four 4-port 100BASE-TX segments)		2.2.1

Table A-2. Model 5000BH Switch Modules

Part Number	Description	Special Notes	Minimum Software Requirement
CL3104007	Model 5625 24-port 10/100 EtherSpeed autosensing switch module		3.2.1
CL3104008	Model 5424FX 24-port 100BASE-FX SMF/MMF EtherSpeed switch module		4.0.1
Model 5000BH TokenSpeed Switch Modules			
CL1704002	Model 5525M 8-port UTP/STP TokenSpeed MCP switch module		2.1.1
CL1704001	Model 5525 8-port UTP/STP TokenSpeed switch module		2.1.1
CL1704003	Model 5524 8-port TokenSpeed switch module (Fiber 802.5j)		2.2
CL1704004	Model 5525HD 24-port TokenSpeed switch module		3.2.1
Model 5000BH ATMSpeed MCP and Switch Modules and Media Dependent Adapters			
CL1304016	Model 57622M-MM ATMSpeed MCP switch module		3.2.1
CL1304015	Model 57622M-SM ATMSpeed MCP switch module		3.2.1
CL1304014	Model 5720M ATMSpeed MDA MCP switch module		2.2.1
CL1304002	Model 5724M 4-port MMF ATMSpeed MCP switch module		2.0.3
CL1304017	Model 57622-MM ATMSpeed switch module		3.2.1
CL1304001	Model 5724 4-port MMF ATMSpeed switch module		2.1.0
CL1304004	Model 5724S 4-port SMF ATMSpeed switch module		2.1.0
CL1333001	2-port SONET/SDH MMF ATMSpeed MDA		2.2.1
CL1333002	2-port SONET/SDH UTP ATMSpeed MDA		2.2.1
CL1333003	2-port SONET/SDH SMF ATMSpeed MDA		2.2.1
CL1333004	2-port DS-3 ATM MDA		3.2.1
CL1333005	2-port E-3 ATM MDA		3.2.1

Switch Image Compatibility

Table A-3 lists the available MCP modules and the release versions that support them. For upgrade information, refer to *Release Notes for the Centillion Platform*.

Table A-3. Switch Image Compatibility Matrix

MCP Type	2.1	2.1.1	2.2	3.0.1 Std	3.0.1 Adv	3.1.1 Std	3.1.1 Adv	3.2 Std	3.2 Adv	4.0 Std	4.0 Adv
4-port TokenSpeed/ MCP	Yes	No	No	No	No	No	No	No	No	No	No
8-port TokenSpeed/ MCP	N/A	Yes	Yes	Yes	No	Yes	No	Yes	No	Yes	No
Model 5525M MCP	N/A	Yes	Yes	Yes	No	Yes	No	Yes	No	Yes	No
ATMSpeed/155 MCP	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
ATMSpeed MDA MCP	N/A	N/A	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
ATMSpeed 622 MCP	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	Yes	Yes
Model 5724M MCP	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Model 5720M MDA MCP	N/A	N/A	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Model 57622M MM/SM MCP	N/A	N/A	N/A	N/A	N/A	Yes	Yes	Yes	Yes	Yes	Yes
C50T switch	N/A	Yes	Yes	Yes	No	Yes	No	Yes	No	Yes	No
C50N switch	N/A	Yes	Yes	Yes	No	Yes	No	Yes	No	Yes	No

Caution: You must have a switch image that is compatible with your MCP switch module. Loading an incompatible switch image will corrupt your flash memory.

Caution: The ATMSpeed MCP module must have a minimum of 16 MB of DRAM to support the Advanced software; however, 20 MB is required to support Advanced software if you run LNM.

Index

Numbers

802.1q
- frame tagging 1-13
- VLANs 1-14
- VLANs and MultiLink Trunking 1-20

A

ADD_GROUP_NAME_QUERY frames 7-2

ADD_NAME_QUERY frames 7-2

additional destination filters 6-18

Address Resolution Protocol. *See* ARP

address summarization 4-22

Admin MAC Address, creating 1-56

administrative weight. *See* cost

advanced image 4-3

All Routes Explorer. *See* ARE

alternate-path routing 1-36

ARE
- and Spanning Tree explorer frames, dropping 6-9
- and SRE frames, dropping 6-11
- frames 6-7

ARP and RARP frames, dropping 6-16

ATM
- Address in routing table 5-5
- connectivity 1-40
- distributed switch configuration 1-43
- using LAN emulation 1-41

ATM Forum web site 1-49

ATM routing engine. *See* ARE 3-1

B

Backbone Node 3-1

backbones
- ATM 1-40
- collapsed 2-9, 2-18
- increasing bandwidth on
  - Ethernet 2-15, 2-18
  - token ring 2-4, 2-9
- replacing 2-18

Backplane Connection field 1-33

backplane ring, connecting or disconnecting 1-33

Banyan frame filters 6-17

BN 3-1

border nodes 4-3, 4-22

bridges, replacing 2-4, 2-15

broadcast and unknown server. *See* BUS

BUS
- redundant services 1-37

BUS, definition 1-3

C

Call 5-4

call routing and load balancing 5-12

Centillion LAN Client. *See* CLC

Centillion Multiprotocol Engine (MPE) 3-1

Circuit Saver
- compared to Turbo 1-44
- full mesh connectivity 1-47
- switching concepts and examples 1-44
- with CLC 1-4
- with LEC 1-5

Circuit Saver forwarding method 1-4
CLC
virtual ports 1-5
with Circuit Saver 1-4
CLC Turbo LAN. *See* GIGArray LAN
CN_BN filter 6-7
collapsed backbone 2-9, 2-18
common SAPs 6-27
compatibility with ATM 1-40
complete prefix in routing table 5-5
concentrator ports and segmentation 2-2
configuration
concepts 1-26
examples 2-1
Configure Call Route dialog box 5-4
connectivity,
using ATM 1-40
control ATM address 3-4
core groups 4-5
core switch 4-10
cost value in routing table 5-6
cost, redundancy with IISP 1-37

D

data ATM address 3-4
DATAGRAM frames 7-2
dedicated port for a server 2-2, 2-13
distributed backbone switching
with LANE 2-12
without LANE 2-10
DTL
multilevel 4-21
dual paths, preventing 1-30

E

edge switches 4-12
educational services xix
ELAN, definition 1-2
Emulated LAN. *See* ELAN
Ethernet 802.3 6-4
Ethernet LAN, segmenting 2-13
Ethernet portion ATM MCP module 1-26
Ethernet ports
dropping frames 6-10, 6-15
forwarding 802.3 SNA frames 6-12
forwarding version 2.1 SNA frames 6-13
monitoring frames 6-19
Ethernet protocol types 6-28
Ethernet V2.0 6-4
external rings 1-53

F

FDDI, compared with ATM 1-40
filters
definition 1-21, 6-5
examples 6-6
outbound examples 6-26
frame tagging 1-13
frame types, NetBIOS 7-2
Front Port Connect field 1-34

G

Get Nearest Server broadcasts (IPX)
local ring 6-22
workgroup rings 6-23
GIGArray LAN
compared to Circuit Saver LAN 1-44
example configuration 1-48
switching concepts and examples 1-46

H

hop count
- effect of dedicated server port 2-2
- reducing 2-8

HP Laser Jet filters, reducing broadcasts 6-25

I

ID numbers, definition 1-26

IGMP queries 1-22

IGMP/IP 1-22

IISP
- alternate route 5-9
- configuring links 5-3
- definition 1-6
- link cost 5-6
- link groups 5-4, 5-6, 5-7
- link groups (example) 5-13
- link states 4-13
- operational requirements 5-3
- preventing loops 5-7
- providing redundancy 1-36
- redundant links (example) 5-11
- route table 5-3
- routing information entry 5-5
- routing tables 5-3
- with PNNI 4-9

Interim Inter-switch Signaling Protocol. *See* IISP

internal rings 1-53

IP filters
- dropping Ethernet 2.0 frames (except ARP and RARP) 6-15
- dropping token ring ARP and RARP frames 6-16
- dropping token ring frames (except ARP and RARP) 6-15

IP Multicast 1-22

IPX filters
- dropping Ethernet IPX traffic 6-14
- dropping token ring IPX traffic 6-13
- Get Nearest Server broadcasts
  - restricting to local ring 6-22
  - restricting to workgroup rings 6-23

IRAP packets 1-23

L

LAN Emulation Client. *See* LEC

LAN Emulation Configuration Server. *See* LECS

LAN Emulation Server. *See* LES

LAN Emulation. *See* LANE

LAN Network Manager. *See* LNM

LAN performance, improving
- concentrating LANs into a router 2-5, 2-16
- dedicating ports to a server 2-2, 2-13
- Ethernet 2-13
- replacing backbones 2-18
- replacing bridges and routers 2-4, 2-15
- segmenting a token ring 2-2

LANE
- definition 1-1
- virtual ports 1-5

LANE services
- implementation considerations 1-49
- implementing 1-49
- recommendations 1-51

LANs
- concentrating into a router 2-5, 2-16
- connecting using ATM 1-40

latency, decreasing 1-40

LEC
- definition 1-2
- virtual ports 1-5
- with Circuit Saver 1-5

LECS
- definition 1-3
- redundancy 1-50
- redundant services 1-37

LES
- definition 1-3
- redundant services 1-37

LES/BUS pairs 1-37

LES/BUS, implementation considerations 1-50

limits
- MLT 1-18

link cost 5-6

link groups 5-4, 5-6, 5-11
LLC offset 6-3
LNM 1-52
load balancing and sharing
  IISP 5-12
  LEC Circuit Saver and Turbo virtual ports 1-36
  LECS 1-38
  LES/BUS pairs 1-37
  Turbo 1-47
local ring
  and IPX Get Nearest Server broadcasts 6-22
  restricting RIP and SAP broadcasts on 6-24
logical links 4-19
logical nodes 1-7, 4-22
loop detection 5-6
LUNI 2.0. *See* LANE.

M

MAC offset 6-3, 6-4
Management Wrap field 1-32
master control processor. *See* MCP
MCP
  limits 1-51
  memory requirements 1-52
  NetBIOS processing 7-1
  redundancy 1-39
  with multiple bridge groups 1-14
memory requirements, MCP 1-52
Model 5000BH switch
  distributed backbone switching 2-10
Model 5525HD switch module 2-11
module number in routing table 5-5
monitor filters
  Ethernet ports 6-19
  token ring ports 6-18
MPC
  definition 1-10, 3-3
  egress functionality 3-5
  ingress functionality 3-4
MPE 3-1
MPOA
  components 3-3
  definition 1-10, 3-1
  discovery 3-6
  establishing a shortcut 3-7
  internal shortcut 3-10
  MPC 3-3
  MPS 3-5
  server 3-5
  TLV 3-6
  VCC multiplexing 3-10
MPOA client. *See* MPC
MPOA server. *See* MPS
MPS
  definition 3-5
  egress functionality 3-6
  ingress functionality 3-6
MPS, definition 1-10
multilevel DTL 1-7, 4-21
MultiLink Trunking
  and 802.1q VLANs 1-20
  definition 1-18
  limits 1-18
multiplexing VCCs 3-10
Multi-Protocol Over ATM. *See* MPOA.

N

name filters, NetBIOS 7-2
name proxy frames, NetBIOS 7-2
NAME_QUERY frames 7-2
NetBIOS
  filtering 1-21
  frame types 7-2
  internal filter 7-1
  name filters 7-3
  name filters (examples) 7-2
  name proxy 7-2, 7-6
  name table 7-6
  processing through the MCP 7-1
NetBIOS filters
  dropping frames from ports 6-10
  dropping SRE and ARE frames 6-11

Network Service Access Point. *See* NSAP
Network-to-Network Interface. *See* NNI
Next Hop Resolution Protocol. *See* NHRP
NHRP
  and MPOA 3-9
  definition 1-10
NNI, definition 1-6
nonzero virtual ports, signaling 1-35
nonzero VPIs, for Circuit Saver virtual ports 1-5
NSAP
  addresses 4-14
NSAP, addresses 4-2
NSAPs, registered 4-14

O

outbound filters 6-26

P

packet filters
  and HP Laser Jet broadcasts 6-25
  Banyan frames 6-17
  CN_BN 6-7
  common SAPs 6-27
  directing routed protocols 6-20
  general explorer filters 6-9
  introduction 6-3
  IP filters 6-15
  IPX filters 6-13
  limitations 6-1
  managing traffic 6-2
  monitor filters 6-18
  NetBIOS filters 6-10
  offsets 6-3, 6-4
  outbound 6-26
  overview 6-1
  restricting IPX Get Nearest Server broadcasts 6-22
  restricting RIP and SAP broadcasts 6-24
  sequence limitations 6-5
  SNA filters 6-12
partial prefix in routing table 5-5
path balancing 4-25
peer group
  definition 4-2
  ID 4-17
  leaders 4-3
performance, Circuit Saver compared with Turbo 1-44
physical ports within VLANs, definition of 1-15
PNNI
  and QoS 4-2
  configuring 4-14
  core designs 4-7
  definition 1-7, 4-1
  dialog box 4-17
  features 4-18
  hierarchy 4-3
  multilevel DTL 1-7, 4-21
  network designs 4-6
  path balancing 4-25
  peer groups 4-2
  routing protocol 4-1
  signaling protocol 4-2
  SpeedView dialog box 4-17
  with IISP 4-9
  with SpeedView 4-3
port number in routing table 5-5
port-based VLAN 1-12
ports, dedicating to a server 2-2
private NNI 1-6
Private-Network-to-Network Interface. *See* PNNI
product support xix
protocol-based VLAN 1-12
Pseudo query mode 1-22
publications
  Bay Networks xviii
  related xviii
PVC encapsulation method 1-8

R

RARP frames, dropping 6-16

redundancy
- and IISP 5-11
- Centillion 50/100 and Model 5000BH networks 1-36
- LECS 1-38, 1-50
- LES/BUS pairs 1-37, 1-50
- spanning tree 1-36
- Turbo 1-47
- with ATM MCPs 1-39
- with Circuit Saver and Turbo virtual ports 1-39
- with LANE virtual ports 1-36

redundant MCP 1-39

related publications xviii

Reverse Address Resolution Protocol. *See* RARP

Ring In/Ring Out ports and segmentation 2-2

ring numbers 1-56

ring parameter server. *See* RPS

RIP filters 6-24

route types 5-5

routed protocols, directing 6-20

routers
- adding memory and processing cycles 2-6
- concentrating rings into 2-5, 2-16
- defining subnets 2-6
- relieving of source route bridging 2-6
- replacing 2-4, 2-15

routing tables
- PNNI 4-9
- setting up 5-3

RPS 1-56

S

SAP filters 6-24

segmentation of a LAN 2-2, 2-13

server, improving performance 2-2, 2-13

service access points (SAPs) 6-27

shortcut, establishing an MPOA 3-8

Single Route Explorer. *See* SRE

SNA filters
- forwarding Ethernet 802.3 SNA frames 6-12
- forwarding Ethernet version 2.0 SNA frames 6-13
- forwarding token ring SNA frames 6-12

soft PVCs, definition 1-8

SONET, description 1-9

source route bridging, removing from a router 2-6

source route hop count 2-2

source route, replacing multiple two-port bridges 2-4

Spanning Tree explorer frames, dropping 6-9

spanning tree groups 1-12

Spanning Tree Protocol, redundancy 1-36

SRE frames, dropping 6-11

star topology 1-42

STATUS_QUERY frames 7-2

subnets, defining 2-6

support, Bay Networks xix

SVCs (Switched Virtual Circuits)
- and PNNI 4-3
- redundancy 1-36

switch
- displaying with LNM 1-55

Switch Connect field 1-31

Synchronous Digital Hierarchy. *See* SDH

Synchronous Optical Network. *See* SONET

T

targetless ARP 1-11

technical publications xviii

technical support xix

TLV (type length value) 3-4, 3-6

token ring filters 6-20

token ring LAN, segmenting 2-2

token ring ports
- dropping frames 6-15
- dropping SRE and ARE frames 6-11
- forwarding SNA frames 6-12
- monitor filters 6-18

token ring-to-ATM connectivity 1-40

Turbo
- CLC 1-44
- forwarding method 1-4
- load sharing 1-47
- redundancy 1-47
- virtual ports 1-5

type length value. *See* TLV

U

UNI, definition 1-6

User-to-Network Interface. *See* UNI

V

VCC multiplexing 3-10

VCI, for Circuit Saver virtual ports 1-5

Virtual LAN. *See* VLAN

virtual path connection. *See* VPC

virtual ports
- CLC 1-5
- configuring LANE ports for redundancy 1-36
- definition 1-4
- LEC 1-5
- load balancing 1-39
- managing with LNM 1-55
- nonzero 1-35
- redundancy in a network 1-39
- signaling over nonzero 1-35
- with Turbo 1-4
- within VLANs, definition 1-16

virtual rings
- definition 1-18
- reducing source route hops 2-8

virtual routing 3-2

VLAN
- 802.1q 1-14
- configuring 1-17
- creating 2-6
- definition 1-11
- in spanning tree groups 1-12
- multiple 1-16
- port-based 1-12
- protocol-based 1-12
- with physical ports 1-15
- with virtual ports 1-16

VLAN ID 1-14

VPC 4-19

W

workgroup rings
- and IPX Get Nearest Server broadcasts 6-23
- restricting RIP and SAP broadcasts 6-24